

625-CD-610-003

EOSDIS Core System Project

ECS Training Material Volume 10: Archive Processing

July 2003

Raytheon Company
Upper Marlboro, Maryland

ECS Project Training Material Volume 10: Archive Processing

July 2003

Prepared Under Contract NAS5-60000
CDRL Item 129

RESPONSIBLE ENGINEER

<u>Paul E. Van Hemel /s/</u>	<u>7/24/2003</u>
Paul E. Van Hemel	Date
EOSDIS Core System Project	

SUBMITTED BY

<u>Gary. Sloan /s/</u>	<u>7/24/2003</u>
Gary Sloan, M&O Manager	Date
EOSDIS Core System Project	

Raytheon Company
Upper Marlboro, Maryland

This page intentionally left blank.

Preface

This document is a contract deliverable with an approval code of 3. As such, it does not require formal Government approval. This document is delivered for information only, but is subject to approval as meeting contractual requirements.

Any questions should be addressed to:

Data Management Office
The ECS Project Office
Raytheon Company
1616 McCormick Dr.
Upper Marlboro, MD 20774-5301

This page intentionally left blank.

Abstract

This is Volume 10 of a series of lessons containing the training material for the Earth Observing System Data and Information System (EOSDIS) Core System (ECS). This lesson provides a detailed description of the process required to perform the tasks associated with archive functions.

Keywords: training, archive, AMASS, ACSLS, AAWin, granule deletion tool, Spatial Subscription Server, Data Pool, Data Pool maintenance, course objective

This page intentionally left blank.

Change Information Page

List of Effective Pages			
Page Number		Issue	
Title		Revised	
iii through xii		Revised	
1 through 243		Revised	
Slide Presentation 1 through 99		Revised	
Document History			
Document Number	Status/Issue	Publication Date	CCR Number
625-CD-610-001	Original	March 2001	
625-CD-610-002	Revised	September 2002	
625-CD-610-003	Revised	July 2003	

This page intentionally left blank.

Contents

Preface

Abstract

Introduction

Identification	1
Scope.....	1
Purpose.....	1
Status and Schedule	1
Organization.....	1

Related Documentation

Parent Document.....	3
Applicable Documents.....	3
Information Documents	3
Information Documents Referenced.....	3
Information Documents Not Referenced.....	3

Archive Processing

Lesson Overview	7
Lesson Objectives	7
Importance	11

Overview of Archive Processing

Hardware	12
Software	15
ACSLs	16
Data Sources and Uses.....	17

Starting and Shutting Down AMASS

Starting the AMASS Tape Archive System	20
Shutting Down AMASS Tape Archive System.....	22
Rebooting AMASS	23
Entering the Archive After AMASS is Started.....	24

Archive Storage Structures

Storage Element Relationships	26
Launching DSS GUIs	28
Archive Resource Management.....	32

Insert Data Into the Archive

Archive Insert Scenario	36
-------------------------------	----

Monitor Archive Requests

Request Status Window	40
Distrib'n Requests Window	44

Retrieve Data From The Archive

Monitoring Distribution Requests	46
--	----

Deleting Granules

Deletion Capability and Features.....	48
Deletion Sequence	49
Undelete Capability	54

Loading Archive Media

Automatically Loading Archive Media	55
Manually Loading Archive Media.....	56
Formatting a Volume	57
Remove Media	58

Backup Archived Data

Creating Offsite Backups	60
Creating a Backup for AMASS	63
Replacing the AMASS Database Backup Volume (Volume 1)	64
Create Replacement Backups Manually from Existing Archives	66

Restore Archive Data

Use of Backup Data for Recovery	67
Manual Data Recovery from Local or Offsite Backup Tapes	67
Reprocessing	68
Requesting Replacement Data from Provider	68
Restoring the AMASS Database	81
Restoring the ACSLS Database	82

AMASS Graphical User Interface

Modify a Volume Group	86
Modify a Volume	89

Archive Monitoring and Troubleshooting

AMASS Commands, Utilities, and Scripts for Monitoring and Fault Response	96
Recovery from Failure to Store or Retrieve Data	105
Checking Server Log Files	105
A Special Case: Checking the Request Manager Server Debug Log	112
Checking the <i>tac</i> Log	123
Handling a Data Insertion Failure	125
Handling a Data Acquire Failure	129
Checksum De-activation	134

Data Pool Management

Data Pool Architecture.....	137
Using the Data Pool Maintenance GUI	140
Monitoring Data Pool Active Insert Processes and Insert Actions.....	158
Checking the Status of Batch Inserts	159
Suspending and Resuming Data Pool Insert Actions	159
Checking the Data Pool Insert Queue and Canceling a Data Pool Insert Action	160
Working with the NoFreeSpaceFlag.....	162
Managing Data Pool Configuration Parameters and Data Pool Tuning.....	162
Managing Data Pool Collection Groups.....	166
Managing Data Pool Collections within Collection Groups	169
Using the DPM GUI to Manage Themes.....	174
Troubleshooting DPM GUI Problems	179
Using the Spatial Subscription Server (NSBRV) GUI	180
Working with Existing Subscriptions.....	191
Using the NSBRV GUI to Manage Bundling Orders.....	194
Using the NSBRV GUI for Monitoring the Spatial Subscription Server	200
Working with Data Pool Scripts	201
Extending the Period of Retention for Granules in the Data Pool.....	202
Running the Data Pool Cleanup Utility	205
Running the Data Pool Access Statistics Utility.....	211
Using the Batch Insert Utility	220
Working with the DataPool Order Status & Control GUI.....	222
Orders Page.....	222
Order Items Page.....	224
Queue Control Page	226

Practical Exercises

Perform Activities Related to Archive Processing	234
--	-----

Slide Presentation

Slide Presentation Description.....	236
-------------------------------------	-----

Introduction

Identification

Training Material Volume 10 is part of Contract Data Requirements List (CDRL) Item 129, whose requirements are specified in Data Item Description (DID) 625/OP3 and is a required deliverable under the Earth Observing System Data and Information System (EOSDIS) Core System (ECS), Contract (NAS5-60000).

Scope

Training Material Volume 10 describes the process and procedures associated with Archive Processing. It describes archive hardware, software, and data. In addition, it addresses starting and shutting down the tape archive control software, monitoring archive requests, and performing archive management tasks. This lesson is designed to provide the operations staff with sufficient knowledge and information to satisfy all lesson objectives.

Purpose

The purpose of this Student Guide is to provide a detailed course of instruction that forms the basis for understanding data archiving. Lesson objectives are developed and will be used to guide the flow of instruction for this lesson. The lesson objectives will serve as the basis for verifying that all lesson topics are contained within this Student Guide and slide presentation material.

Status and Schedule

This lesson module provides detailed information about training for the current baseline of the system. Revisions are submitted as needed.

Organization

This document is organized as follows:

Introduction:	The Introduction presents the document identification, scope, purpose, and organization.
Related Documentation:	Related Documentation identifies parent, applicable and information documents associated with this document.
Student Guide:	The Student Guide identifies the core elements of this lesson. All Lesson Objectives and associated topics are included.
Slide Presentation:	Slide Presentation is reserved for all slides used by the instructor during the presentation of this lesson.

This page intentionally left blank.

Related Documentation

Parent Document

The parent document is the document from which this ECS Training Material's scope and content are derived.

423-41-01	Goddard Space Flight Center, EOSDIS Core System (ECS) Statement of Work
-----------	---

Applicable Documents

The following documents are referenced within this ECS Training Material, or are directly applicable, or contain policies or other directive matters that are binding upon the content of this document:

420-05-03	Goddard Space Flight Center, Earth Observing System (EOS) Performance Assurance Requirements for the EOSDIS Core System (ECS)
423-41-02	Goddard Space Flight Center, Functional and Performance Requirements Specification for the Earth Observing System Data and Information System (EOSDIS) Core System (ECS)

Information Documents

Information Documents Referenced

The following documents are referenced herein and amplify or clarify the information presented in this document. These documents are not binding on the content of the ECS Training Material.

609-CD-610	Release 6B Operations Tools Manual for the ECS Project
611-CD-610	Mission Operation Procedures for the ECS Project
910-TDA-022	Custom Code Configuration Parameters for ECS

Information Documents Not Referenced

The following documents, although not referenced herein and/or not directly applicable, do amplify or clarify the information presented in this document. These documents are not binding on the content of the ECS Training Material.

305-CD-610	Release 6B Segment/Design Specification for the ECS Project
311-CD-620	Release 6B Data Management Subsystem Database Design and Schema Specifications for the ECS Project
311-CD-621	Release 6B INGEST (INS) Database Design and Schema Specifications for the ECS Project

311-CD-623	Release 6B Planning and Data Processing Subsystem Database Design and Schema Specifications for the ECS Project
311-CD-624	Release 6B Science Data Server Database Design and Schema Specifications for the ECS Project
311-CD-625	Release 6B Storage Management and Data Distribution Subsystem Database Design and Database Schema Specifications for the ECS Project
311-CD-626	Release 6B Subscription Server Database Design and Schema Specifications for the ECS Project
311-CD-627	Release 6B Systems Management Subsystem Database Design and Schema Specifications for the ECS Project
311-CD-628	Release 6B Registry Database Design and Schema Specifications for the ECS Project
311-CD-630	Release 6B PDS Subsystem Database Design and Database Schema Specifications for the ECS Project
311-CD-631	Release 6B NameServer Database Design and Schema Specifications for the ECS Project
313-CD-610	Release 6B ECS Internal Interface Control Document for the ECS Project
334-CD-610	6B Science System Release Plan for the ECS Project
601-CD-001	Maintenance and Operations Management Plan for the ECS Project
603-CD-003	ECS Operational Readiness Plan for Release 2.0
604-CD-001	Operations Concept for the ECS Project: Part 1-- ECS Overview
604-CD-002	Operations Concept for the ECS Project: Part 2B -- ECS Release B
605-CD-002	Release B SDPS/CSMS Operations Scenarios for the ECS Project
607-CD-001	ECS Maintenance and Operations Position Descriptions
152-TP-001	ACRONYMS for the EOSDIS Core System (ECS) Project
152-TP-003	Glossary of Terms for the EOSDIS Core System (ECS) Project
211-TP-007	Transition Plan 6A.04 to 6A.XX (6A.05) for the ECS Project
220-TP-001	Operations Scenarios - ECS Release B.0 Impacts
300-TP-002	Database Descriptions for Synergy III
500-1002	Goddard Space Flight Center, Network and Mission Operations Support (NMOS) Certification Program, 1/90

535-TIP-CPT-001

Goddard Space Flight Center, Mission Operations and Data Systems
Directorate (MO&DSD) Technical Information Program Networks
Technical Training Facility, Contractor-Provided Training
Specification

This page intentionally left blank.

Archive Processing

Lesson Overview

This lesson reviews the process of archiving data, including a description of processing for working with the File Storage Management System (FSMS) software, monitoring the ingest/archiving/distribution performance, managing archive content and capacity, maintaining configuration of peripherals and data servers, backing up and restoring archived data, documenting and troubleshooting archive errors, maintaining the archive processing queue (storage and retrieval), and providing archive status.

Lesson Objectives

Overall Objective - The overall objective of this lesson is proficiency in the methodology and procedures for archive processing in the Earth Observing System Data and Information System (EOSDIS) Core System (ECS) during maintenance and operations. The lesson includes a description of processing for monitoring the ingest/archival/distribution performance, maintaining configuration of peripherals and data servers, documenting archive errors, maintaining the archive processing queue (both storing and retrieval), managing archive content and capacity, submitting new data archive requests to the Science Coordinator, and providing archive status.

Specific Objective 1 - The student will list DAAC operator positions for Archive Manager personnel interfaces and identify responsibilities associated with each interface.

Condition - The student will be given a list of DAAC operators.

Standard - The student will select four personnel positions with which the Archive Manager interfaces and list at least one major area of responsibility for each selected position.

Specific Objective 2 - The student will identify the major hardware facility for archival storage and its associated storage cartridges.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, a working system archive facility, a 9940 tape cartridge, and a 9840 tape cartridge.

Standard - The student will correctly point out the StorageTek Library Storage Module (LSM) and its associated 9940 and 9840 tape cartridges.

Specific Objective 3 - The student will describe the File Storage Management System (FSMS) software.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project* and a copy of 609-CD-610-003 *Release 6B Operations Tools Manual*.

Standard - The student will identify the FSMS software as the Archival Management and Storage System (AMASS), correctly describe AMASS by stating its nature as a UNIX file system installed on an SGI XL computer, and state the five steps in the AMASS control path without error.

Specific Objective 4 - The student will start the AMASS tape archive system.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive facility.

Standard - The student will operate the STK control panels, power up the archive hardware, and then boot AMASS host and start AMASS without error and in accordance with documented procedures.

Specific Objective 5 - The student will shut down the AMASS tape archive system.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive facility with AMASS started.

Standard - The student will terminate AMASS and shut down the LSM without error and in accordance with documented procedures.

Specific Objective 6 - The student will use manual mode to enter the LSM.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive facility with AMASS started.

Standard - The student will vary the LSM offline, enter the LSM, leave the LSM, and vary the LSM back online, without error, in accordance with documented procedures, and following all required safety precautions.

Specific Objective 7 - The student will describe the relationships between Earth Science Data Types (ESDTs), Logical Volume Groups (LVGs) in the Archive, and physical archive volume groups.

Condition - The student will be given a diagram depicting the relationships.

Standard - The student will correctly explain the logical and physical structure of ECS archive storage.

Specific Objective 8 - The student will describe the process of, and monitor the progress of, inserting new data into the archive.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive facility.

Standard - The student will describe without error the process of inserting new data into the archive.

Specific Objective 9 - The student will launch available Data Server Subsystem (DSS) Graphical User Interfaces (GUIs) and monitor retrieval of data from the archive.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive facility.

Standard - The student will perform without error and in accordance with documented procedures the procedures relating to monitoring retrieval of data from the archive.

Specific Objective 10 - The student will use the granule deletion capability to delete granules from the archive and inventory.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive facility.

Standard - The student will perform without error and in accordance with documented procedures the procedure for deleting granules from the archive and inventory.

Specific Objective 11 - The student will perform automatic and manual loading of archive storage cartridges.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, a working system archive facility, 9940 tape cartridges, and 9840 tape cartridges.

Standard - The student will perform procedures for automatic and manual loading of the proper cartridges for the LSM without error and in accordance with documented procedures.

Specific Objective 12 - The student will create a backup for AMASS.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive facility.

Standard - The student will correctly and in accordance with documented procedures use the `vgexport -q` command to create the AMASS backup.

Specific Objective 13 - The student will replace a full Backup Volume.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive facility.

Standard - The student will perform without error and in accordance with documented procedures the procedures for replacing the Backup Volume (Volume 1).

Specific Objective 14 - The student will manually create a replacement backup for an archive data tape.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive facility.

Standard - The student will perform without error and in accordance with documented procedures the procedures for manually creating a replacement backup for an archive data tape.

Specific Objective 15 - The student will “restore” archive data by inserting a backup copy cartridge.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive facility.

Standard - The student will perform without error the procedure for inserting a backup copy cartridge to replace a lost archive data tape.

Specific Objective 16 - The student will conduct basic checks on AMASS status and functioning appropriate for initial troubleshooting.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive facility.

Standard - The student will successfully check for active processes reflecting required AMASS daemons, run commands to check the basic health of the running AMASS applications, and run an AMASS command to display current AMASS input/output activity without error and in accordance with documented procedures.

Specific Objective 17 - The student will use the *quedisplay* command to display what is in the AMASS queue.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive facility.

Standard - The student will perform without error and in accordance with documented procedures the procedure for viewing what is in the AMASS queue.

Specific Objective 18 - The student will use the *amass_log* script to display AMASS errors.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive facility.

Standard - The student will perform without error and in accordance with documented procedures the procedure for using the *amass_log* script to display AMASS messages from the system log file.

Specific Objective 19 - The student will launch and use the AMASS Graphical User Interface (GUI).

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive facility.

Standard - The student will successfully start the AMASS GUI, display icons and data for volume groups and volumes, and execute procedures for modifying volume groups and volumes without error and in accordance with documented procedures.

Specific Objective 20 - The student will perform the functions required to maintain the Data Pool, including tasks with the Data Pool Maintenance (DPM) GUI (monitor Data Pool active insert processes; monitor/cancel data pool insert actions; suspend and resume Data Pool actions; check the Data Pool insert queue; toggle the state of the NoFreeSpace flag; configure the number of allowed active insert processes; configure the default retention period and the default retention priority; view and update collection groups in the Data Pool database; list/add/delete a theme), tasks with utilities and scripts (extend the period of retention for selected science granules already in the Data Pool; set up a schedule and cron job for Data Pool cleanup; manually invoke Data Pool cleanup; set up a schedule and cron job for Data Pool access statistics accumulation; manually invoke the Data Pool access statistics utility), and tasks with the Spatial Subscription Server (NSBRV) GUI (extend the period of retention in a Data Pool insert subscription; update a subscription; view/add/cancel a bundling order; view statistics on processing of events and actions by the Spatial Subscription Server).

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working Data Pool facility.

Standard - The student will use the GUI tools, scripts, and utilities without error in accordance with applicable procedures to perform the required Data Pool maintenance functions.

Specific Objective 21 - The student will launch and use the DataPool Order Status & Control GUI to view Data Pool orders and order items, and to control the HEG Converter Front End Server.

Condition - The student will be given a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, 609-CD-610-003 *Release 6B Operations Tools Manual*, and a working system archive and Data Pool facility.

Standard - The student will successfully start the Order Status & Control GUI, display a list of orders, review order items and details, and execute procedures for starting and stopping the HEG Converter Front End Server and adjusting its parameters without error and in accordance with documented procedures.

Importance

The Archive Manager's role in maintaining the archive data is key to the successful implementation and operation of ECS. Ensuring the smooth operation of the archive is crucial for ECS core functionality.

Overview of Archive Processing

Archive processing is at the heart of the Earth Observing System Data and Information System (EOSDIS) Core System (ECS) at the Distributed Active Archive Centers (DAACs). Through archive processing, data that have been ingested into the system are archived to tape for permanent storage and distributed to users via hard media (tape or disk) or electronic means.

The DAAC Archive Manager's job entails working with the Science Data Specialist, the Science Coordinator, and the Resource Manager, as well as providing direction for the Data Ingest Technician. These personnel interfaces are illustrated in Figure 1.

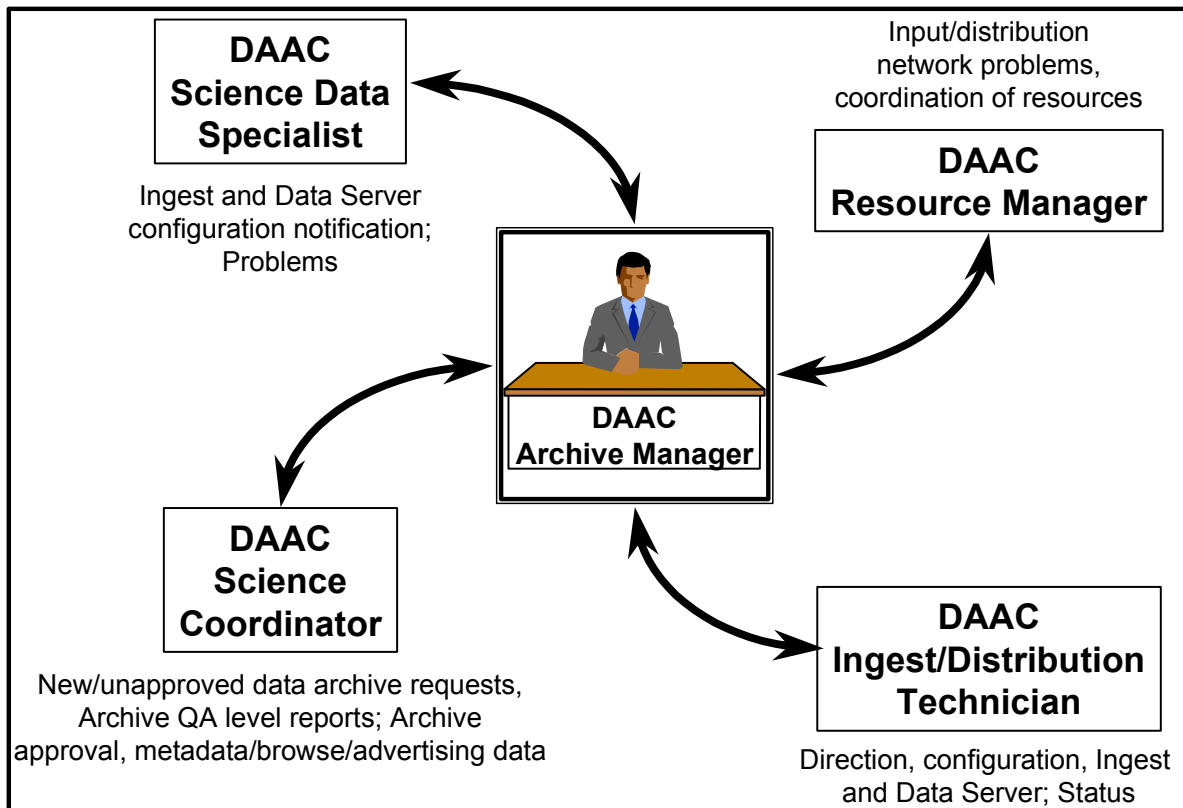


Figure 1. Archive Manager Personnel Interfaces

The Data Server Subsystem (DSS) is responsible for persistent storage of earth science and related data, for providing search and retrieval access to this data, and for supporting the administration of the data, hardware devices, and software products. As part of its retrieval function, the subsystem also provides for distribution of data electronically or on physical media.

Hardware

The ECS Archive uses one major type of archival storage hardware for storing science data, browse data, and other ECS data. The StorageTek (STK) Powderhorn Model 9310 Automated

Cartridge System tape storage tower, illustrated in Figure 2, is a mass storage system of removable media jukeboxes. The software that manages the storage in the ECS architecture is hosted on a Silicon Graphics Inc. (SGI) Origin 2000.

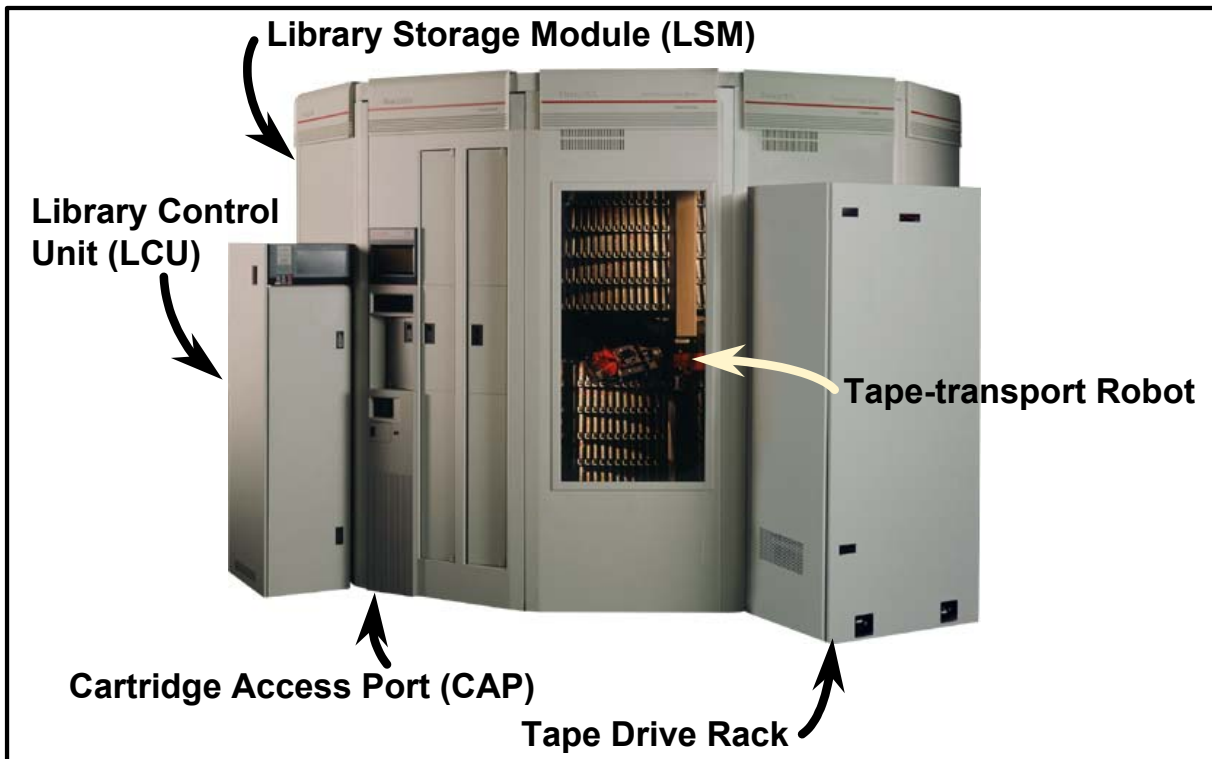


Figure 2. StorageTek (STK) Powderhorn (Model 9310) Data Storage

The typical ECS data storage archive consists of the following major elements:

- Library Storage Module (LSM), Powderhorn Model 9310, Automated Cartridge System (ACS) tape storage tower.
- Cartridge Access Port (CAP), where media are inserted or ejected from the LSM; standard capacity is 21 cartridges.
- dual tape-transport "robots" for moving cartridges from the tower to a tape drive or CAP and from the tape drive or CAP to the tower.
- tape drive rack with eight 9940 cartridge tape drives (rack capacity is 20 drives).
- Library Management Unit (LMU), Model 9330, a serial port for the ACS Library Software (ACSLs) that controls and monitors the ACS.
- Library Control Unit (LCU), Model 9311, a hardware interface for managing LSM intercommunications.
- browse tape drive rack with eight 9840 tape drives (rack capacity is 20 drives).

The LSM tape archive can store thousands of tapes. The archive stores science data on STK 9940 tapes, each of which can store 60 gigabytes of data (up to 200 gigabytes compressed). Browse data are stored on 9840 tapes, each capable of storing 80 gigabytes (compressed) of data. Each 9940 tape cartridge is identified by a colored bar code label that shows the media number (see Figure 3). The 9840 tapes are of the same physical dimensions and use bar code labels, as shown in Figure 4. An archive catalog or database tracks the location of each cartridge within the library, based on information provided by the laser bar code reader.

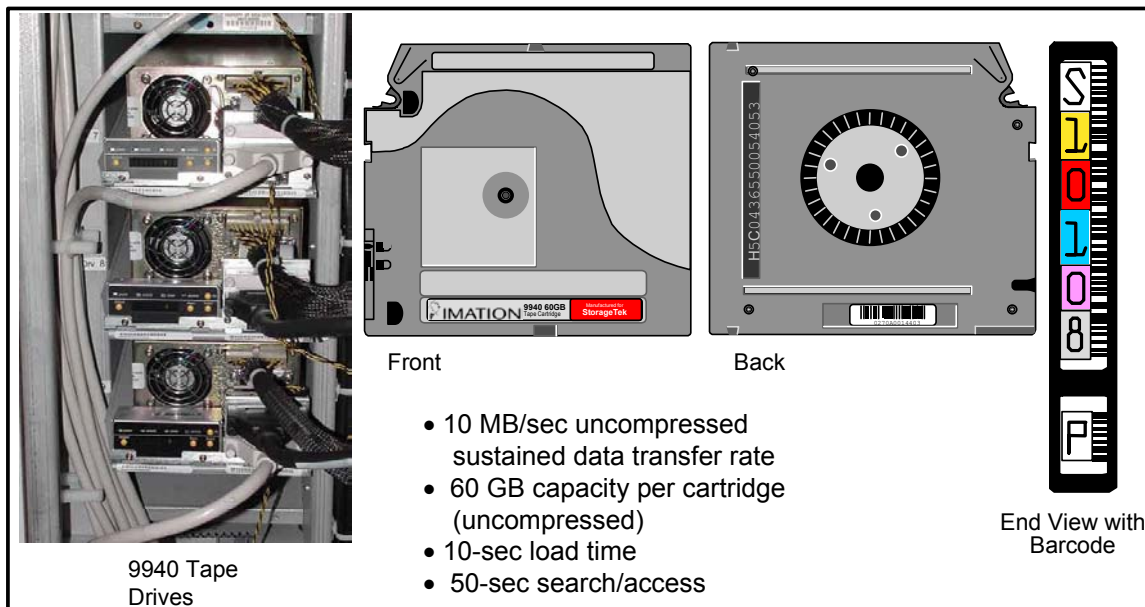


Figure 3. STK 9940 Tape System

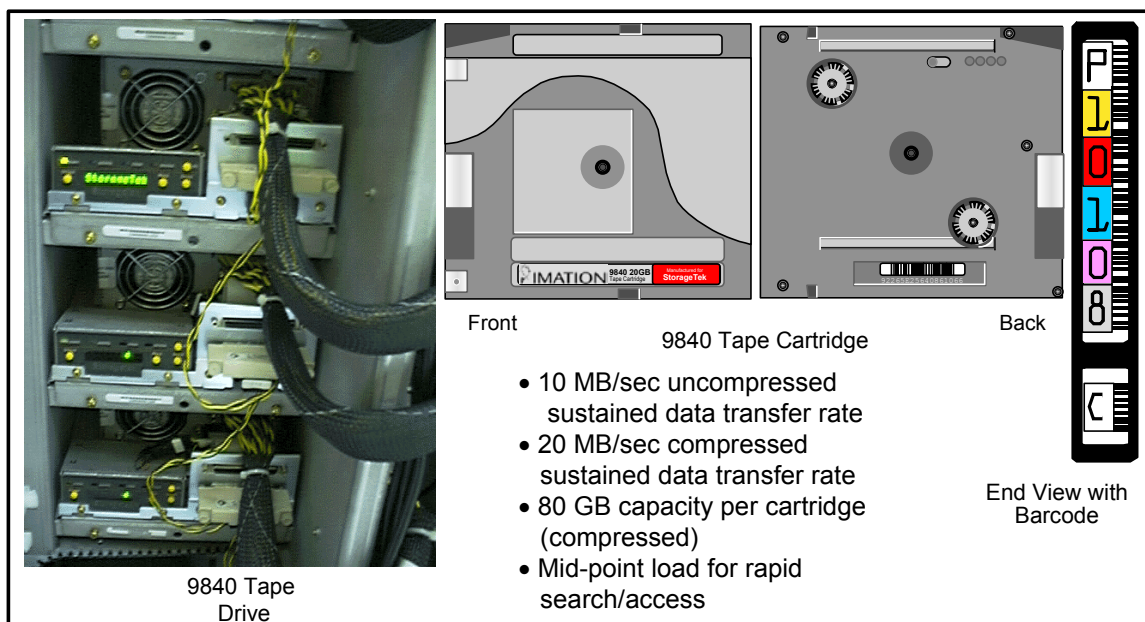


Figure 4. 9840 Tape System for Browse Data

Software

Archive operations rely on both custom and commercial off the shelf (COTS) software for complete mass storage archive management, providing the capability to accept Graphical User Interface (GUI) and command line interface inputs, and to interpret them to the appropriate level needed to control and monitor archive operations. The File Storage Management Service (FSMS) software is the Archival Management and Storage System (AMASS), a product of Advanced Digital Information Corporation (ADIC). The purpose of AMASS in the ECS is to provide an easy-to-use interface to a large tape archive. AMASS is a UNIX file system that manages files, volumes (media), drives and jukeboxes. It allows UNIX File System (UFS) access methods to be employed (e.g., ftp, rcp, uucp, nfs, RPC, native) while removing some of the limitations of the UFS. Primary among these is reliance on UNIX Index Node (inode) structures. AMASS maintains all inode information in database files rather than in associated disk structures. This minimizes or eliminates many of the file search problems inherent in searching large numbers of files in multiple directories. In addition, AMASS organizes files as groups of blocks which can be individually retrieved. This differs from UFS resident systems that require staging the entire file.

The ECS software for the archive provides a standardized format and content for logging of data access and staging activity, which may assist in system troubleshooting. The system also provides parallel AMASS input/output capability for optimum system throughput. Furthermore, the system uses a logical archive ID capability for complete separation of the physical location of data in the archive from its logical reference in the inventory. This means that client requests for data do not change based on a change of the physical location of the data, and as a result there can be improved load balancing, cross-archive fault recovery, and archive upgrades.

AMASS is installed on an SGI Origin 2000 computer. Control information is communicated from the SGI to the LMU using TCP/IP protocols via the FDDI network. Figure 5 shows the basic route taken by control information in the process of sending a file to AMASS.

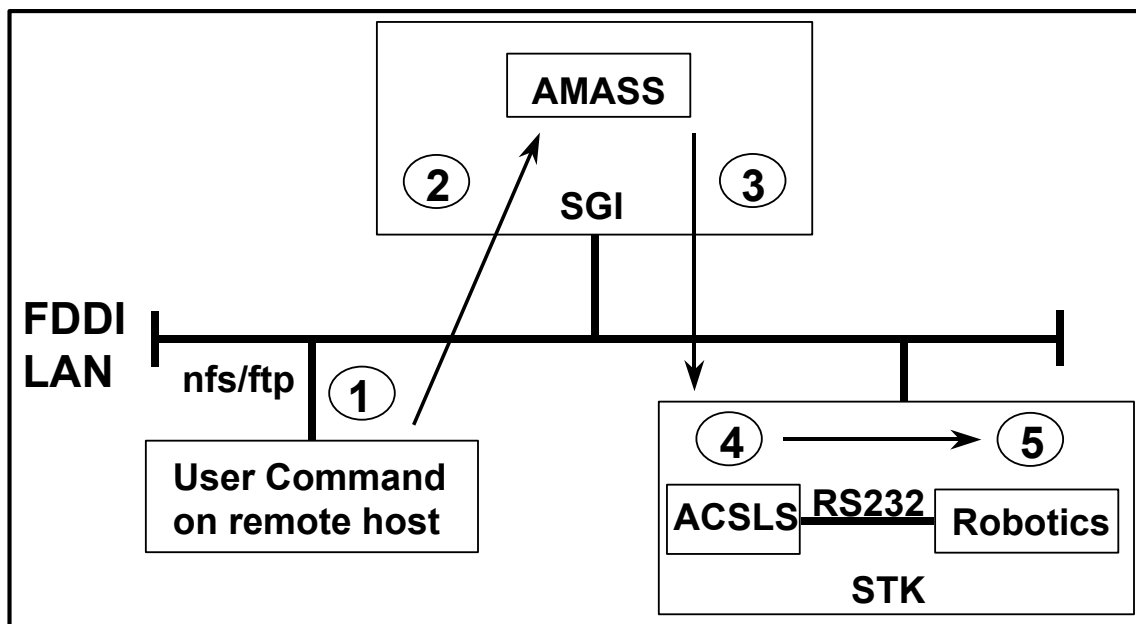


Figure 5. AMASS Control Path

As the figure suggests, there are five elements in the path:

1. The user or application initiates transfer of the file to AMASS.
2. AMASS receives the file over the network via nfs, ftp, dd, or cp, or locally via dd or cp to its cache.
3. AMASS sends information to ACSLS to specify the tape to load.
4. The ACSLS application sends the LMU robot to retrieve the tape.
5. The robot grips the tape, retrieves it, and inserts it into the tape drive to complete the mount.

ACSLS

For the StorageTek Powderhorn, Steps 4 and 5 in the AMASS control path are managed by the Automated Cartridge System Library Software (ACSLS), running on a Sun computer under Solaris 2.6. Full guidance for using ACSLS is provided in the *Automated Cartridge System Library Software System Administrator's Guide, Version 5.3*. Table 1 lists the commands covered in that *Guide*.

Table 1. ACSLS Command Reference

Command	Function
audit	Creates or updates the database inventory of the volumes in a library component.
cancel	Cancels a current or pending request.
clear lock	Removes all active and pending locks on transports or volumes
define pool	Creates or modifies scratch pools.
delete pool	Deletes empty scratch pools.
dismount	Dismounts a volume.
eject	Ejects one or more volumes from the Automated Cartridge System (ACS).
enter	Sets a Cartridge Access Port (CAP) to enter mode.
idle	Stops ACSLS from processing new requests.
lock	Locks (dedicates) a volume or transport to a user.
logoff	Exits the command processor.
mount	Mounts a data or scratch volume.
query	Displays the status of a library component.
set	Sets various attributes of different library components.
show	Displays your lock ID or user ID.
start	Starts ACSLS request processing.
unlock	Removes active locks on volumes or transports.
vary	Changes the state of an ACS, LSM, CAP, transport, or port.
venter	Enters one or more volumes with missing or unreadable labels into the ACS.

ACSLS commands use the following general syntax:

command type_identifier state [options]

where **type_identifier** is the ACS component and its identifier (these are listed in the *System Administrator's Guide*), **state** is a device state for the **vary** command only, and **options** are command options (these are specified for each command in the *System Administrator's Guide*). The two most useful commands in ACSLS are **query** and **vary**. Other frequently used commands are **enter** and **eject**, for inserting and removing cartridges, respectively. ACSLS does not have an online help facility, but if you enter a command (e.g., **vary**), it will prompt you for the parameters.

There are also several utilities provided with ACSLS. These are listed in Table 2.

Table 2. ACSLS Utilities

Utility	Function
bdb.acsss	Backs up the ACSLS database.
kill.acsss	Terminates ACSLS.
rc.acsss	Starts and recovers ACSLS.
rdb.acsss	Restores the ACSLS database.
Volrpt	Creates a volume report.
db_command	Starts or stops the Oracle database.

To control and interact with ACSLS, you use the following user IDs:

- **acssa** lets you enter ACSLS commands from a command processor window.
- **acsss** lets you run ACSLS utilities from the UNIX command line prompt.

It is typical to log in as both user IDs to permit entering both ACSLS utilities and commands. You can, however, open a command processor window from the **acsss** user ID if you prefer to work from a single user ID. The *System Administrator's Guide* provides full details.

Data Sources and Uses

Data that are inserted into the archive are managed by the Data Server Subsystem (DSS) and can be received from such sources as the ingest subsystem, processing subsystem, other DAACs, and authorized users (Figure 6). Uses of data from these sources include:

- from ingest – any ECS function that uses data (e.g., production).
- from processing – various ECS functions (e.g., further processing, distribution to users).
- from other DAACs – various ECS functions (e.g., may be needed as inputs for production of other products).
- from authorized users (via ingest) – typically for distribution or processing.

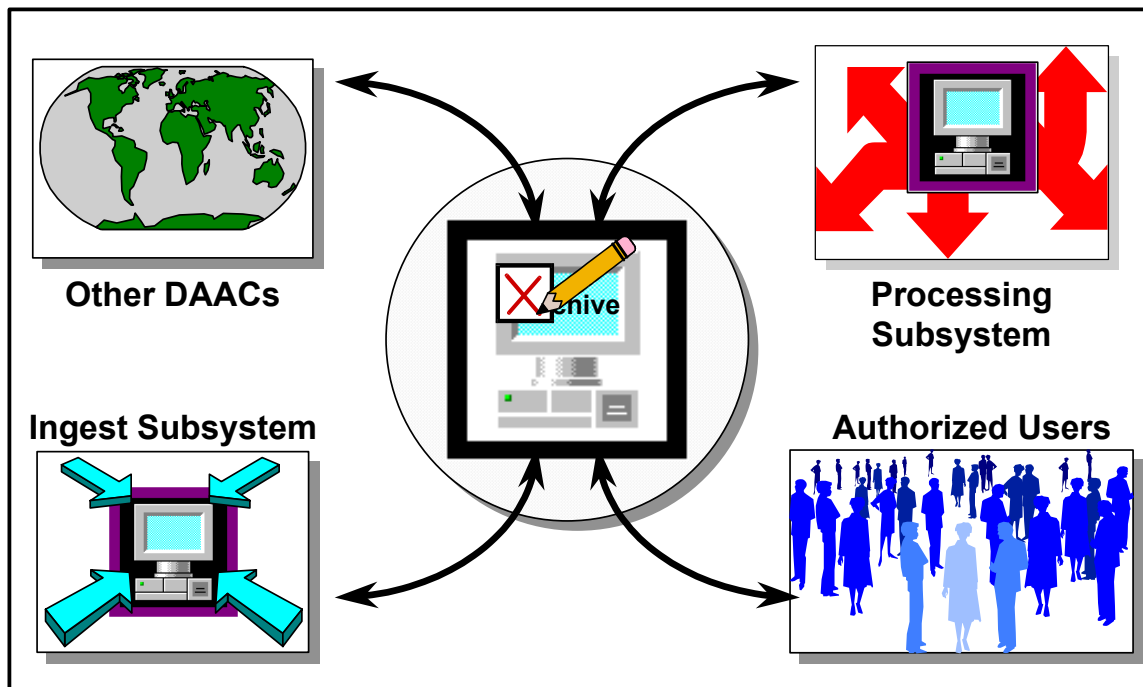


Figure 6. Sources and Uses of Archive Data

This page intentionally left blank.

Starting and Shutting Down AMASS

The AMASS Tape Archive System can be started and shutdown with little or no impact on the rest of the ECS.

Starting the AMASS Tape Archive System

Starting the AMASS FSMS requires actions to ensure that the STK Powderhorn storage system is powered up as well as actions at the SGI FSMS host. Powering up the STK will require actions at its control panels, including the Library Management Unit (LMU) and Library Control Unit (LCU) [the Library Storage Module (LSM) is powered through the LCU]. Figure 7- Figure 9 illustrate major locations of relevant controls and displays.

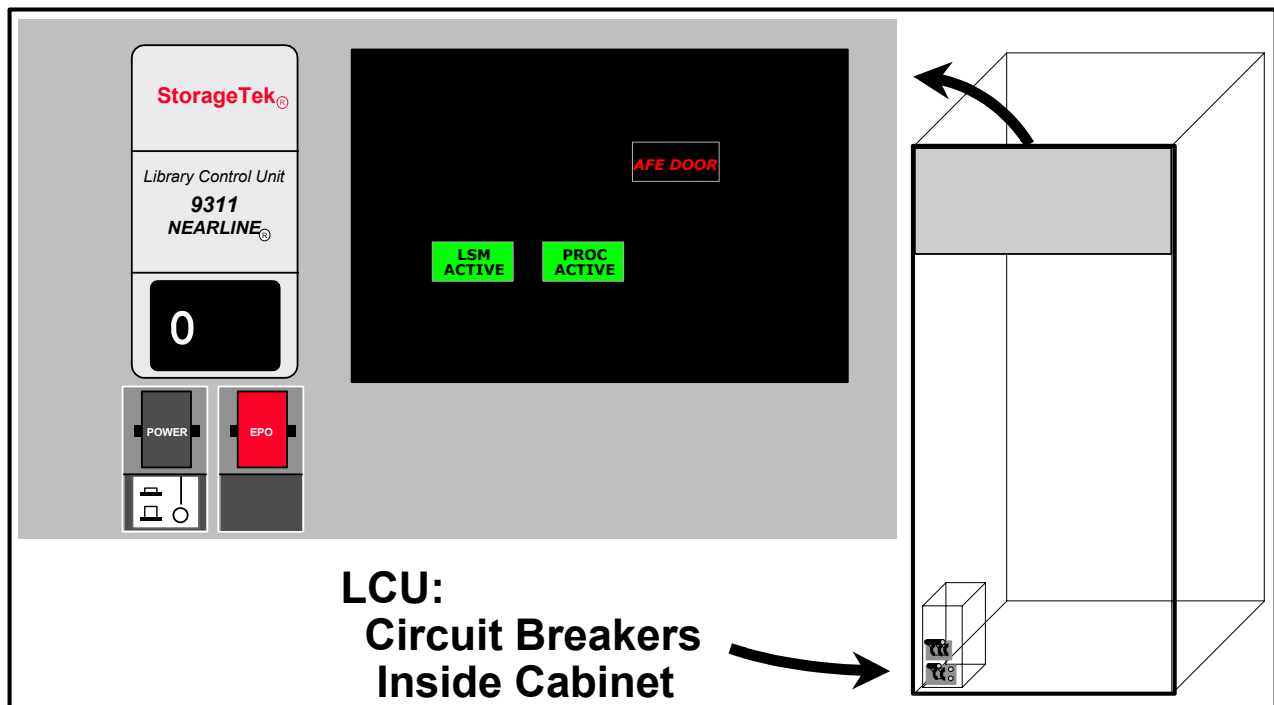


Figure 7. Controls/Displays for the STK Library Control Unit (LCU)

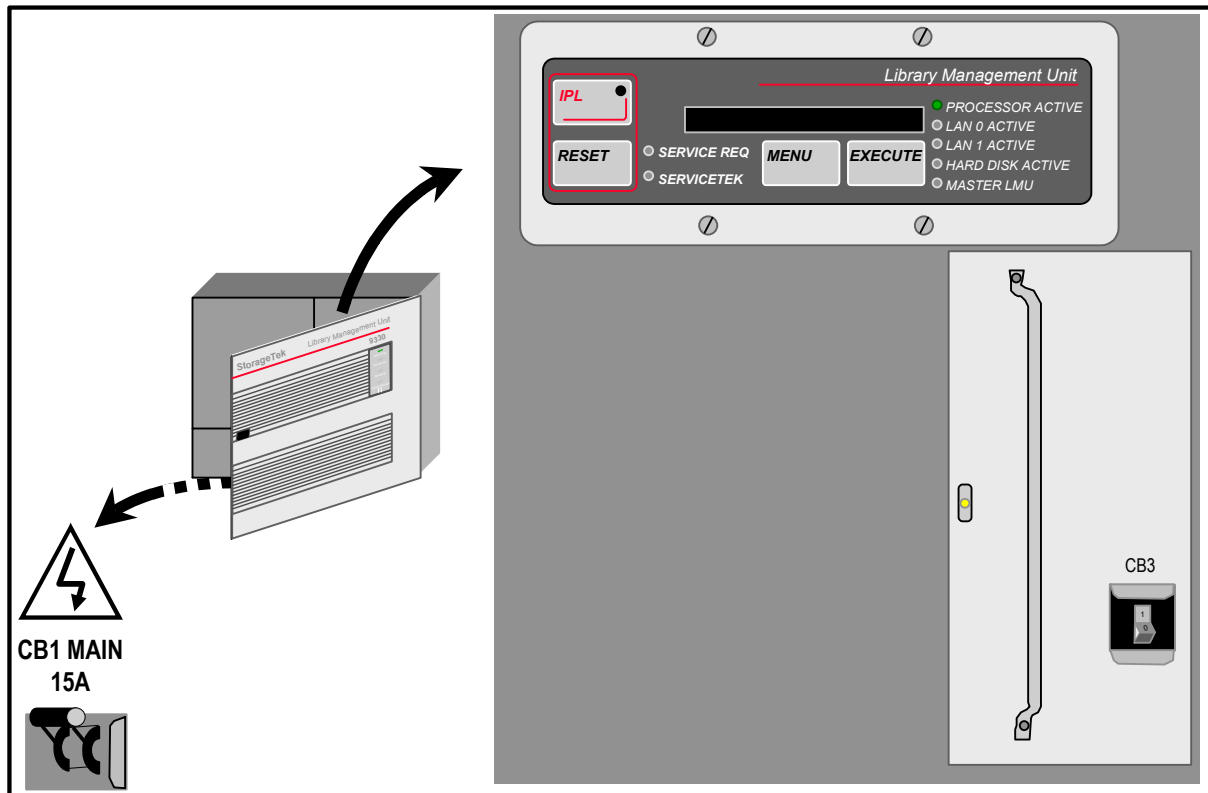


Figure 8. Controls/Displays for the STK Library Management Unit (LMU)



Figure 9. Controls/Displays for the STK Library Storage Module (LSM)

Use the following procedure to start the AMASS software.

Starting the AMASS Application

Note: Preconditions include that 1) the FDDI network is up and running and 2) power to all units is functional and available.

- 1 Make sure power switches for the StorageTek LCU and LMU are **ON**.
 - *NOTE:* The LCU should be the last unit powered up, but otherwise there are no dependencies within the group.
 - 2 If it is not already running, boot the FSMS SGI host normally.
 - There are no dependencies on other hosts, COTS or custom software.
 - AMASS normally starts automatically on bootup. If it does, go to **Step 4**. If it does not, or if you are restarting AMASS after a shutdown, go to **Step 3**.
 - 3 At the FSMS SGI host (workstation **x0drg##**, **xacg##**, or **xwkg##**), as a system administrator (logged in as **root**), type **/usr/amass/tools/amass_start** and then press the **Return/Enter** key.
 - *NOTE:* The **x** in the workstation name will be a letter designating your site: **g** = GSFC, **m** = SMC, **l** = LaRC, **e** = EDC, **n** = NSIDC, **o** = ORNL, **a** = ASF, **j** = JPL; the **##** will be an identifying two-digit number (e.g., **n0drg01** indicates an FSMS SGI server at NSIDC).
 - The AMASS application starts.
 - 4 To verify that AMASS has started correctly, type **/usr/amass/bin/amassstat -c** and then press the **Return/Enter** key.
 - The message **FILESYSTEM IS ACTIVE** is displayed.
-

Shutting Down AMASS Tape Archive System

If it is necessary to shut down AMASS, use the following procedure.

Shut Down the AMASS Application

- 1 Log in as **root** (system administrator) at the FSMS SGI host (workstation **x0drg##**, **xacg##**, or **xwkg##**).
 - *NOTE:* The **x** in the workstation name will be a letter designating your site: **g** = GSFC, **m** = SMC, **l** = LaRC, **e** = EDC, **n** = NSIDC, **o** = ORNL, **a** = ASF, **j** = JPL; the **##** will be an identifying two-digit number (e.g., **n0drg01** indicates an FSMS SGI server at NSIDC).

2 Type **/usr/amass/tools/killdaemons**.

- A message is displayed indicating that all daemons have been terminated.

Rebooting AMASS

The AMASS file system may need to be rebooted during certain anomalous conditions (e.g., system "hang," interruption of communication between AMASS and ACSLS, a required daemon is down). AMASS needs to have the following daemons running at all times: amassmain, daemons/lm_ip -a fslock, klogd, amass_iocomp, qset, libsched, libio_tape,. To verify they are running, simply search for the AMASS processes (refer to Procedure 17.7.1.1 **Checking Daemons and Using *healthcheck*** in Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project*). To check the health of AMASS while it is still running, execute the **healthcheck** command (refer to Procedure 17.7.1.1 in Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project*).

In order to reboot AMASS you must have root privileges. Use the following procedure.

Rebooting AMASS

- 1 Log in as **root** (system administrator) at the FSMS SGI host (workstation **x0drg##**).
 - NOTE: The **x** in the workstation name will be a letter designating your site: **g** = GSFC, **m** = SMC, **l** = LaRC, **e** = EDC, **n** = NSIDC, **o** = ORNL, **a** = ASF, **j** = JPL; the **##** will be an identifying two-digit number (e.g., **n0drg01** indicates an FSMS SGI server at NSIDC).
 - 2 To kill the daemons, type **killdaemons** and then press the **Return/Enter** key.
 - A message is displayed indicating that all daemons have been terminated.
 - 3 If you want to test AMASS before restarting, go to step 4; otherwise, type **amass_start** and then press the **Return/Enter** key.
 - The AMASS application starts.
 - 4 To test the AMASS filesystem prior to starting AMASS type: **install_tests**, and press the **Return/Enter** key.
 - Tests the jukebox operation and cache partitions, then restarts AMASS.
-

Entering the Archive After AMASS is Started

If it is necessary to enter the STK Powderhorn after AMASS is started, use the following procedure.

Entering the STK Powderhorn

- 1 At the host for ACSLS (e.g., **e0drs03**, **g0drs03**, **l0drs02**, **n0drs03**), log in using the **acssa** user ID and password.
 - The **acssa** command-process window is displayed with the **ACSSA>** prompt.
- 2 Type **vary lsm 0,0 offline** and then press the **Return/Enter** key.
 - The access port is unlocked (audible unlatching sound).
- 3 Use the key to unlatch and open the access door.
 - A red **DO NOT ENTER** warning is visible inside the enclosure.

Warning

If it is necessary to enter the STK Powderhorn after AMASS is started, it is necessary to perform the following step to avoid hazard and ensure safety of personnel and equipment.

- 4 Remove the key from the door to ensure that no one inadvertently locks the enclosure with someone inside.
 - The red **DO NOT ENTER** warning is extinguished and a green **ENTER** message is displayed inside the enclosure.
 - 5 Upon leaving the enclosed area, insert the key in the access door and latch the door.
 - The LED display indicates that the door is locked.
 - 6 At the ACSLS host, type **vary lsm 0,0 online** and then press the **Return/Enter** key.
 - After a few seconds, the archive robots execute an initialization sequence and the LSM is back online.
-

This page intentionally left blank.

Archive Storage Structures

Although the physical archive stores data on cartridges (referred to as volumes), it is treated in the system as a large UNIX directory. As such, it is prudent to apply good space management practices in managing the archive just as you would for any computer disk. However, there are other factors that must be considered in archive management, addressing unique system requirements (e.g., on-site backups, off-site backups, desirability of physical consolidation of related data).

Storage Element Relationships

It is important for the Archive Manager to know the relationship between physical storage archives (Library Storage Modules, or LSMs) and the Archive Server software applications at the site. For example, a data repository identified as DRP1 is served by the software application EcDsStArchiveServerDRP1.

Subdivisions within LSMs (e.g., for storage of different data types) are reflected in the Storage Management database, where each Volume Group (a logical group of volumes in the archive) has its own path. As suggested in Figure 10, each path maps to an AMASS volume group, and thus to a physical volume group in the archive.

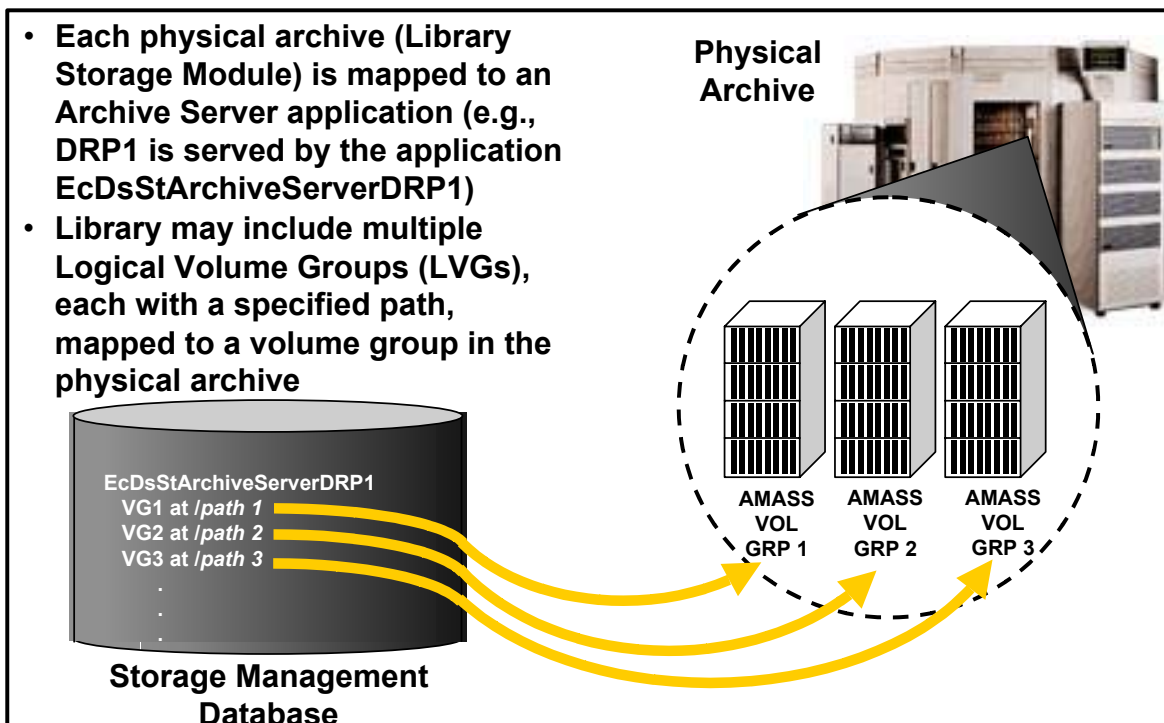


Figure 10. Archive Storage Element Relationships

Information concerning archive servers and the logical volume groups served may be obtained from the Storage Management Control Graphical User Interface (GUI). Specifically, as illustrated in Figure 11, the Storage Configuration tab on the Storage Management GUI permits display of server information and access to related status information.

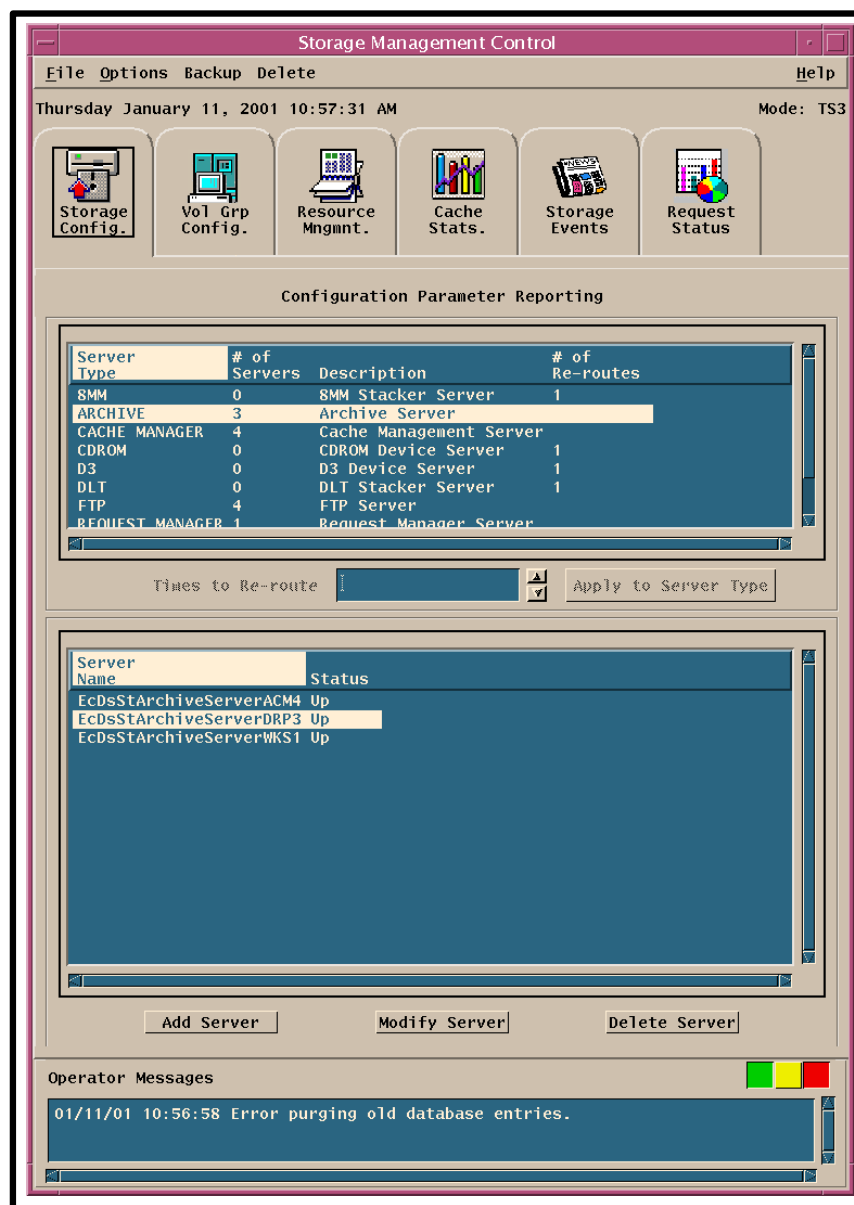


Figure 11. Storage Management, Storage Configuration Tab

The Vol.Grp.,Config. tab, illustrated in Figure 12, permits display of volume group information and history.

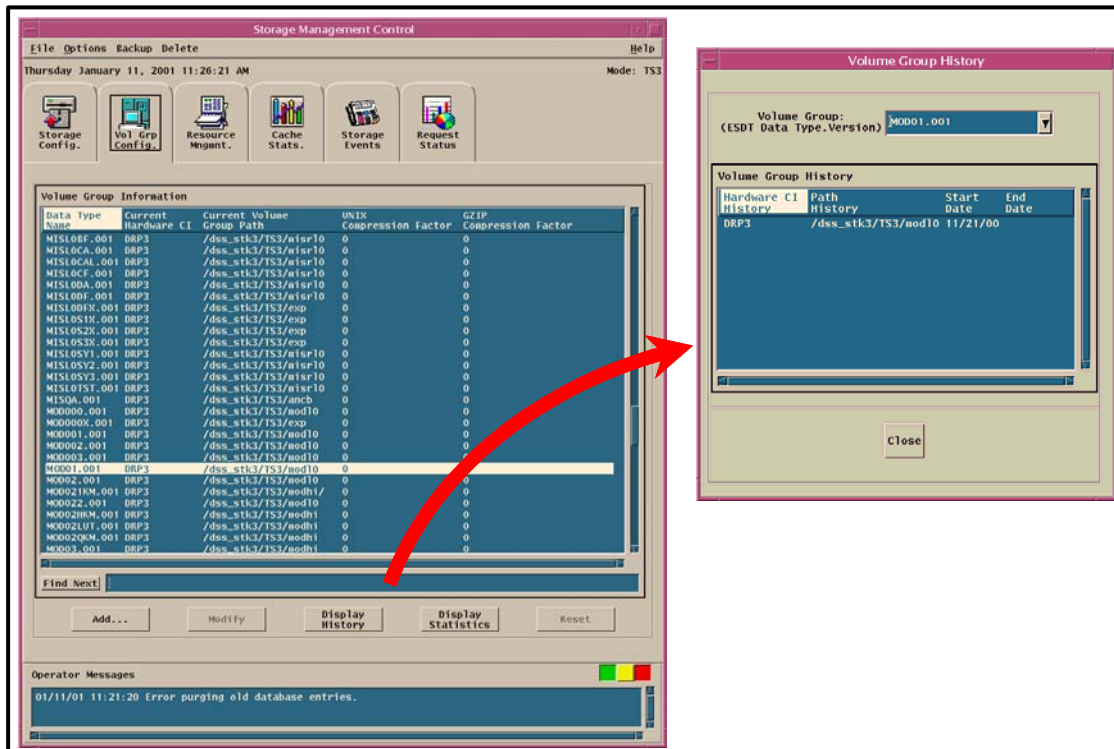


Figure 12. Storage Management, Vol.Grp. Config. Tab

Launching DSS GUIs

The following software applications are associated with DSS:

- Science Data Server (SDSRV).
- Storage Management (STMGT) Servers.
 - Request Manager Server.
 - Staging Disk Server.
 - Cache Manager Server.
 - Archive Server.
 - FTP Server.
 - 9940 Tape Server.
 - 8mm Tape Stacker Server.
- Data Distribution (DDIST) Server.
- DDIST Graphical User Interface (GUI).
- STMGT GUI.
- Science Data Server GUI.

Access to Storage Management, Data Distribution (DDIST), and other GUIs is gained through the use of UNIX commands. The procedure for launching the GUIs begins with the assumption that the applicable servers are running and that the operator (Archive Manager or System Administrator) has logged in.

Launching DSS GUIs Using UNIX Commands

1 Access the command shell.

- The command shell prompt is displayed.

NOTE: Commands in Steps 2 through 9 are typed at a UNIX system prompt.

2 Type **setenv DISPLAY *clientname*:0.0** and then press the **Return/Enter** key.

- Use either the terminal/workstation IP address or the machine-name for the *clientname*.

3 Start the log-in to the appropriate operations workstation (e.g., **e0acs03**, **g0acs02**, **l0acs01**, or **n0acs03**) by typing **/tools/bin/ssh *hostname*** and then press the **Return/Enter** key.

- If you receive the message, **Host key not found from the list of known hosts. Are you sure you want to continue connecting (yes/no)?** type **yes** (“y” alone does not work).
- If you have previously set up a secure shell passphrase and executed **sshremote**, a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears; continue with Step 4.
- If you have not previously set up a secure shell passphrase; go to Step 5.

4 If a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears, type your **Passphrase** and then press the **Return/Enter** key. Go to Step 6.

5 At the **<user@remotehost>'s password:** prompt, type your **Password** and then press the **Return/Enter** key.

6 To change directory to the directory containing the startup scripts for DSS, type **cd /usr/ecs/<MODE>/CUSTOM/utilities** and then press the **Return/Enter** key.

- The **<MODE>** will most likely be one of the following operating modes:
 - OPS (for normal operation).
 - TS1 or TS2 (for testing).
- Note that the separate subdirectories under /usr/ecs apply to different operating modes.

- 7 To launch the Storage Management Control GUI, type the following command: **EcDsStmgtGuiStart <MODE>**, where <MODE> is the one selected in Step 6, and then press the **Return/Enter** key.
 - The Storage Management Control GUI, used for review of storage events and status of devices, is displayed.
 - 8 To launch the Data Distribution GUI, use a similar procedure and type the following command: **EcDsDdistGuiStart <MODE>**, where <MODE> is the one selected in Step 6, and then press the **Return/Enter** key.
 - The Data Distribution GUI is displayed.
 - 9 To launch the **DSS Science Data Server** GUI, use a similar procedure and type the following command: **EcDsSdSrvGuiStart <MODE>** and then press the **Return/Enter** key.
 - The Science Data Server Operator GUI is displayed.
-

Suppose you are an Archive Manager and are asked by someone on the SSI&T team to provide archive path information for a particular Earth Science Data Type (ESDT) stored in the archive. The SSI&T team can tell you the name for the ESDT, because that information is typically in the descriptor file specified when an ESDT is loaded, using the **Data Types** tab of the Science Data Server GUI and the **Add Data Type** dialog illustrated in Figure 13. The figure also illustrates the dialog used to **Update ESDT** information, providing capability permitting updates without reloading ESDTs when changes are necessary.

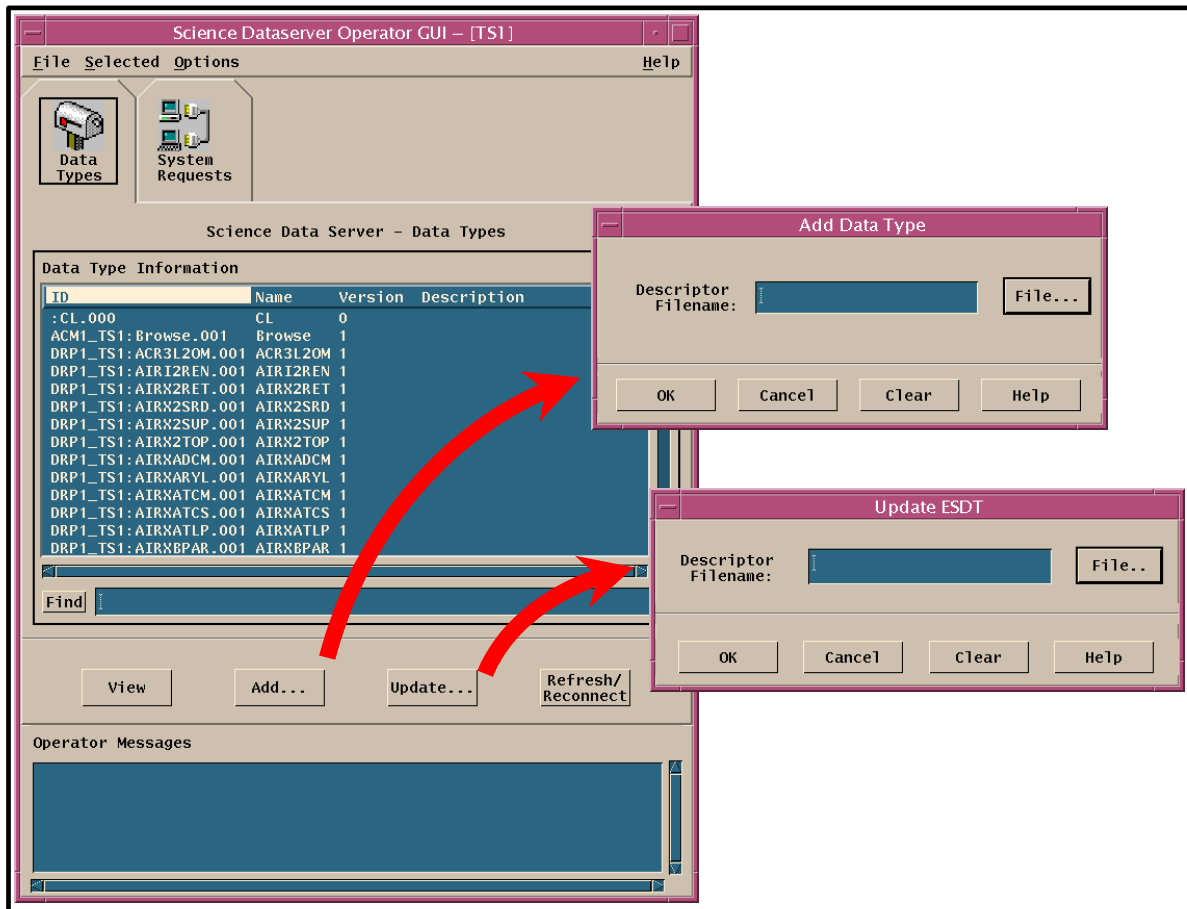


Figure 13. Science Data Server, Data Types Tab and Add/Update Dialogs

Use the following procedure to display specific path information for the archive server.

Use Storage Management GUIs to Display Archive Path Information

- 1 Launch the DSS Storage Management GUI using UNIX commands (see procedure **Launching DSS GUIs Using UNIX Commands** [previous section of this lesson]).
 - The DSS Storage Management GUI is displayed.
- 2 Click on the **Storage Config.** tab to ensure that the Storage Configuration display is available.
 - The **Storage Config.** tab is displayed.
- 3 In the field listing **Server Type**, click on the **ARCHIVE** line to highlight it.
 - The selected line is highlighted and the **Server ID** and **Status** of archive servers are displayed in the field listing **Server Name**.

- 4 Click on the **Vol.Grp. Config.** tab.
 - The **Volume Group Information** is displayed showing volume groups and their current paths.
 - 5 If it is desirable to display the path history for a data type, on the **Vol. Grp. Config.** tab, click on the **Data Type Name** entry for the specific server for which path history information is desired.
 - The selected line is highlighted.
 - 6 Click on the **Display History** button.
 - A **Volume Group History** window is displayed showing the path history for the highlighted data type.
-

Archive Resource Management

The management of archive resources and data is governed by local policy. The software permits sites to establish unique naming and mapping conventions for relating ESDTs to logical volume groups and physical archive volume groups. However, to foster consistency and ease of management of the data in the archive, and to increase the supportability of the system, it is desirable to establish and follow conventions (e.g., naming conventions, assignment of data to volume groups). To facilitate support of local archives by centralized resources (e.g., the System Monitoring and Coordination Center), it is desirable that the DAACs work together (e.g., through Operations Working Groups) to implement consistent and compatible data management practices. Of course, it is necessary for each DAAC to enforce policies and procedures to ensure the long-term viability of archived data.

Logical volume groups are specified using the ESDT short name with the version ID as an extension (e.g., **MOD01.001**). Achieving this convention can be facilitated by adoption of an approach in which specific data products are assigned to the same storage path. In fact, it may be desirable to assign all related products (e.g., **MOD01.001**, **MOD01.002**, **MOD01.003** . . .) to the same storage path. This will have the effect of consolidating related products ultimately in the same physical archive volume group.

The Storage Management GUIs provide tools for managing archive resources. As noted previously, the **Storage Config.** tab (see Figure 11) provides information and control functions for setting and modifying configurations of various Server Types (e.g., 8mm tape, Archive Server) and to manage data location within the archive and on disk. Another Storage Management GUI is the **Resource Mngmnt.** tab, shown in Figure 14. It allows the operator to monitor and adjust the availability of given storage devices. Buttons provide access to pop-up windows permitting management of hardware, including the ability to put specific resources on line or take them off line, and management of media sets.

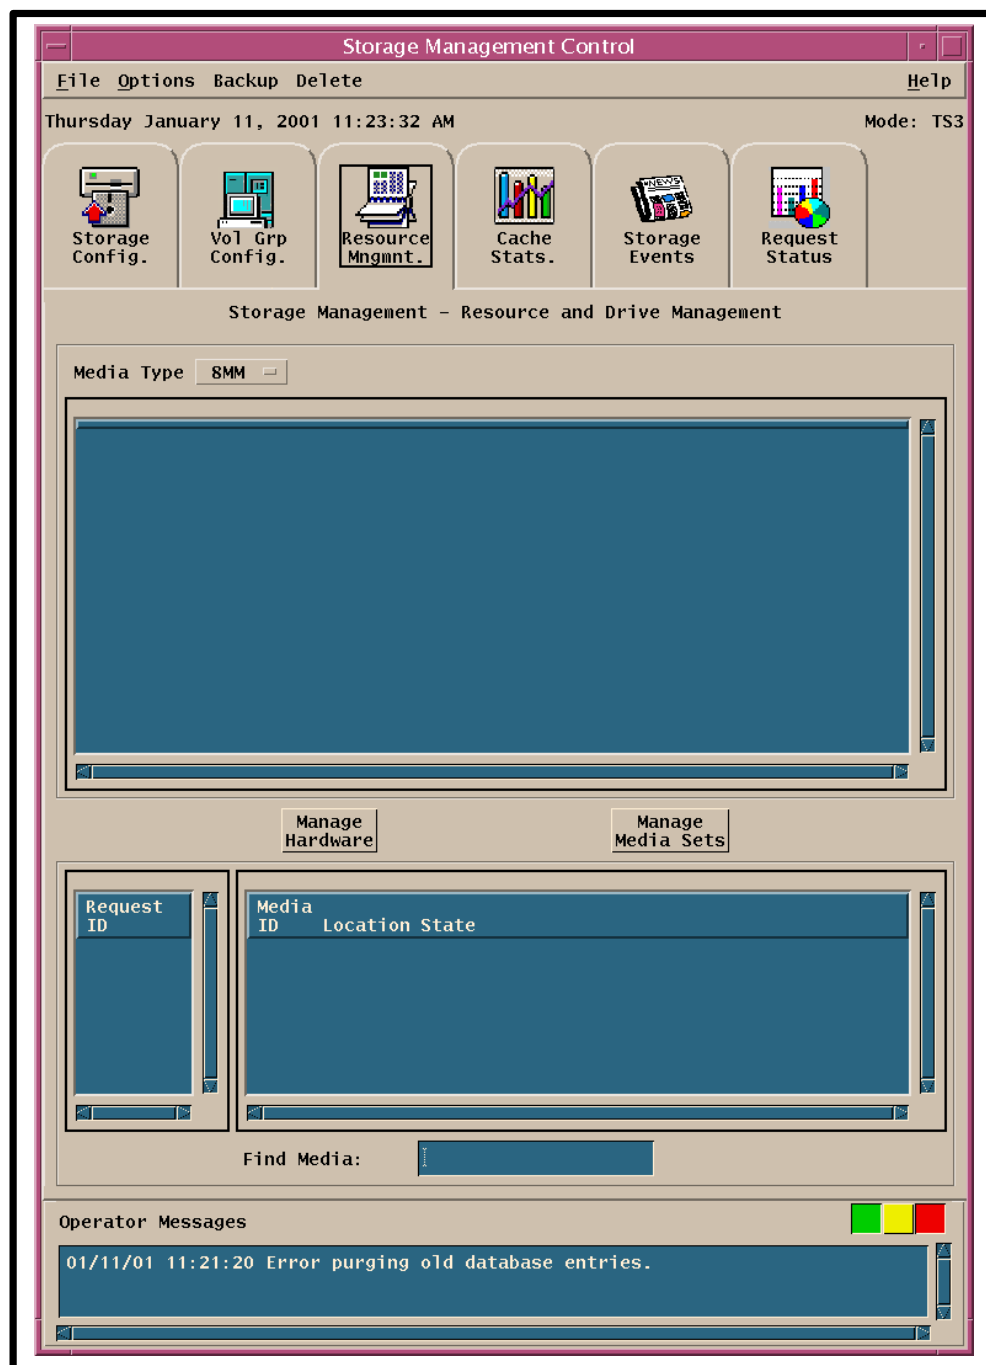


Figure 14. Storage Management, Resource Mngmnt. Tab

Another Storage Management GUI is the **Cache Stats.** tab, shown in Figure 15. It permits the operator to view various system caches, along with statistics on their use.



Figure 15. Storage Management, Cache Stats. Tab

The ECS software implements an active "just-enough" cache cleanup strategy. Under this strategy, the Cache Manager identifies and removes just enough to accommodate new files. Once full, STMGTCaches and the pull area will generally remain full, but this does not require any action by the operator. Instead, the system assesses how much cache or pull disk space is needed and removes expired files only to the extent needed to provide the needed space.

This page intentionally left blank.

Insert Data Into the Archive

Storing new data in the archive repository is largely an automated process that does not normally require operator interaction and occurs as a result of operations such as ingest and data production. Any operations involvement would be to support archive administration operations, resolve problems, periodically monitor working storage and archival operations, and coordinate with the appropriate external/internal sources to resolve schedule conflicts.

Because of the automated nature of this process, step-by-step procedures are not required.

Archive Insert Scenario

NOTE: The scenario that follows describes the insertion of data into the Data Server at an ECS DAAC and is derived from document 605-CD-002-001, *Release B SDPS/CSMS Operations Scenarios for the ECS Project (March 1996)*.

As we have seen, data and associated metadata can be received from numerous sources. This scenario focuses on a routine data insert from the processing subsystem. It assumes that all components are active and not in any degraded modes of operation, that ESDT data collection types have been established, and that the data server's nominal activity rate is 50% of capacity.

Insert Data into the Archive Scenario

- 1 Initiate the session between the Processing Subsystem and the Data Server.
 - The Processing Subsystem sends a Data Insert Request to the Science Data Server.
 - Receipt of the request is logged (via MSS Logging Services) and a request identifier is associated with the Data Insert Request.
 - The content of the request is validated; if successful, it is queued for later processing. If unsuccessful, a rejection message is issued.
 - *The operator may examine the progress of a request by reviewing storage events using the Data Server Subsystem **SDSRV** and **DDIST GUIs** as described in the next section.*

- 2 Transfer data from Processing Subsystem to Data Server.
 - The queued Data Insert Request is reached and processing begins.
 - Associated data granules and metadata are transferred from the Processing Subsystem to the Data Server working storage.
 - Data transfer status, including recoverable errors, is indicated in the event log via MSS Logging Services.
 - *The operator may check the request status using the Data Server Subsystem **SDSRV** & **DDIST GUIs**.*
- 3 Validate metadata received from the Processing Subsystem.
 - The metadata update file(s) produced by the associated product PGEs are validated for completeness and correctness.
 - Validation success or failure is logged via MSS Logging Services with the associated Data Insert Request Identifier and the appropriate status message is returned to the Processing Subsystem.
- 4 Store data granules in the permanent archive.
 - Upon successful validation of the metadata update file, Science Data Server sends a Data Storage Request to Storage Management.
 - The data granules in working storage associated with the Data Storage Request are stored.
 - The Archive Activity Log (via MSS Logging Services) records each data product being stored and storage status of each storage operation.
 - A checksum value is calculated for each data object associated with each granule. (*Note: This calculation can be turned off, and if it is, it may result in the archiving of a corrupted granule with no ready means of detecting the corruption.*)
 - The checksum value (if calculated), storage status, and other selected metadata are forwarded to the Science Data Server in a status message upon completion of the Data Storage Request.

- 5 Store metadata.
- Science Data Server receives and logs the Data Storage Request status message from Storage Management.
 - The additional metadata items are validated.
 - The PGE produced metadata update file and the storage management provided metadata are loaded into the metadata database.
 - The status of the metadata load is entered in the event log.
 - *The operator may examine the progress of a request by reviewing storage events using the Data Server Subsystem **SDSRV** and **DDIST GUIs** as described in the next section.*
- 6 Report Data Insert Request status.
- The Science Data Server logs completion of the Data Insert Request in the event log and reports completion of the Data Insert Request to the Data Archive Manager, the operator console and to the insert Requester (the Processing Subsystem in this scenario).
 - Each of the above entities would also be notified if the request failed and reason(s) for failure identified.
- 7 Process subscriptions based on newly inserted data.
- The Science Data Server will then examine the event list for all subscriptions for that event.
 - Subscription notifications are sent to the appropriate entities as appropriate and distribution processing is initiated.
 - The Science Data Server sends an Advertisement Update Message to the Advertising Server to advertise the new data.
-

This page intentionally left blank.

Monitor Archive Requests

As previously noted, data that are inserted into the archive can be received from a number of sources including the Ingest Subsystem, Processing Subsystem, other DAACs, and Authorized Users. The Data Ingest Technician (DIT) or other operator can monitor the insertion of data into the archive using the Data Server Subsystem (DSS) GUIs.

Request Status Window

A primary GUI tool for monitoring of archive processing is the **Request Status** window (see Figure 16, accessible from the **Storage Management Control** GUI. Using the **Request Status** tab the Archive Manager or Distribution Technician can detect stalled requests or servers that appear to be idle.

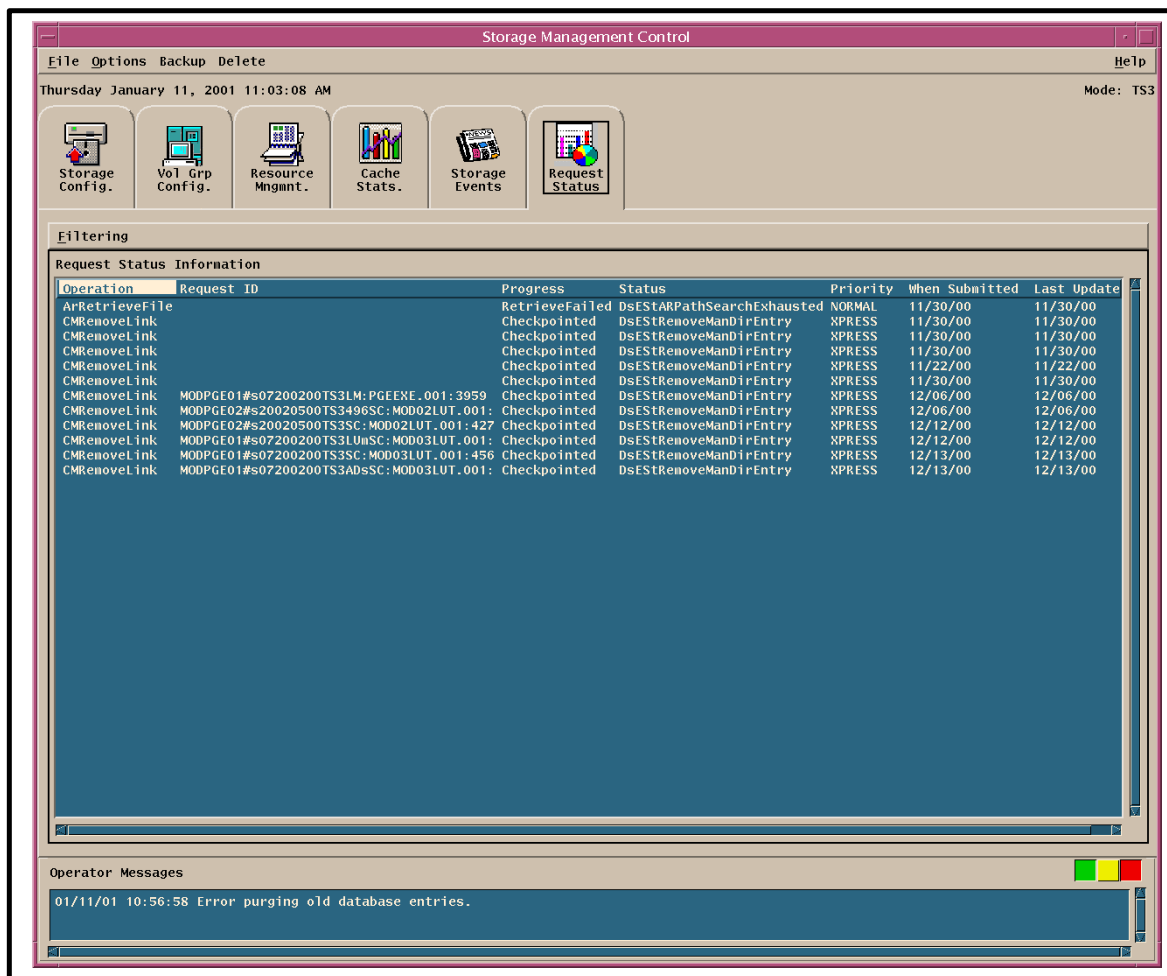


Figure 16. Storage Management Request Status Window

The **Request Status** window displays the following information:

- **Operation** is the type of operation represented by the request.
- **Request ID** is a unique identifier for the request.
- **Progress** is the stage of processing on which the request is currently working (may include a numeric progress indication).
- **Status** provides information about processes attempted and the outcome (e.g., DsEStDRExecuteFailed, DsEStARPathSearchExhausted, OK, . . . WriteFailed, . . .).
- **Priority** is **Xpress**, **Very High**, **High**, **Normal**, or **Low**.
- **When Submitted** is the time and date when the request was received by the Storage Management server that is responsible for the request.
- **Last Updated** is the time and date when the status was last updated for the request.

The operator can reduce the displayed list of requests by clicking on the **Filtering** pull-down menu just above the **Request Status Information** list on the window. This permits filtering on four areas or filter types selectable from the pull-down menu:

- **Server** controls what activity is displayed by limiting the list to the requests being/having been serviced by a specific server. Selecting **All** displays all requests throughout Storage Management. Other selections include the individual archive servers, cache manager servers, ftp servers, request manager server, and staging disk servers.
- **Operation** allows the operator to focus on a specific type of operation. The list of operations is dynamically generated to reflect those operations for which requests are currently in queue (e.g., **All**, **CMLink**, **ArStore**, **ArRetrieve**, **FtpPull**, **FtpPush**).
- **Processing State** allows the operator to differentiate among requests that are being actively processed; have been completed, either successfully or to a retryable error state; or have been suspended and are awaiting the outcome of another event. The following selections are available: **All**, **Processing**, **Suspended**, **Completed**.
- **Submitter** allows the operator to see the status of requests submitted by a specific client process. The list of possible clients is dynamically generated to reflect the list of clients with outstanding requests (e.g., **All**, **DSDD**, **HDFC**, **SDSV**, **this**, [various servers]).

Figure 17 illustrates the Request Filtering menus. The following procedure may be used to monitor archive requests using the Storage Management GUI and Request Status window.

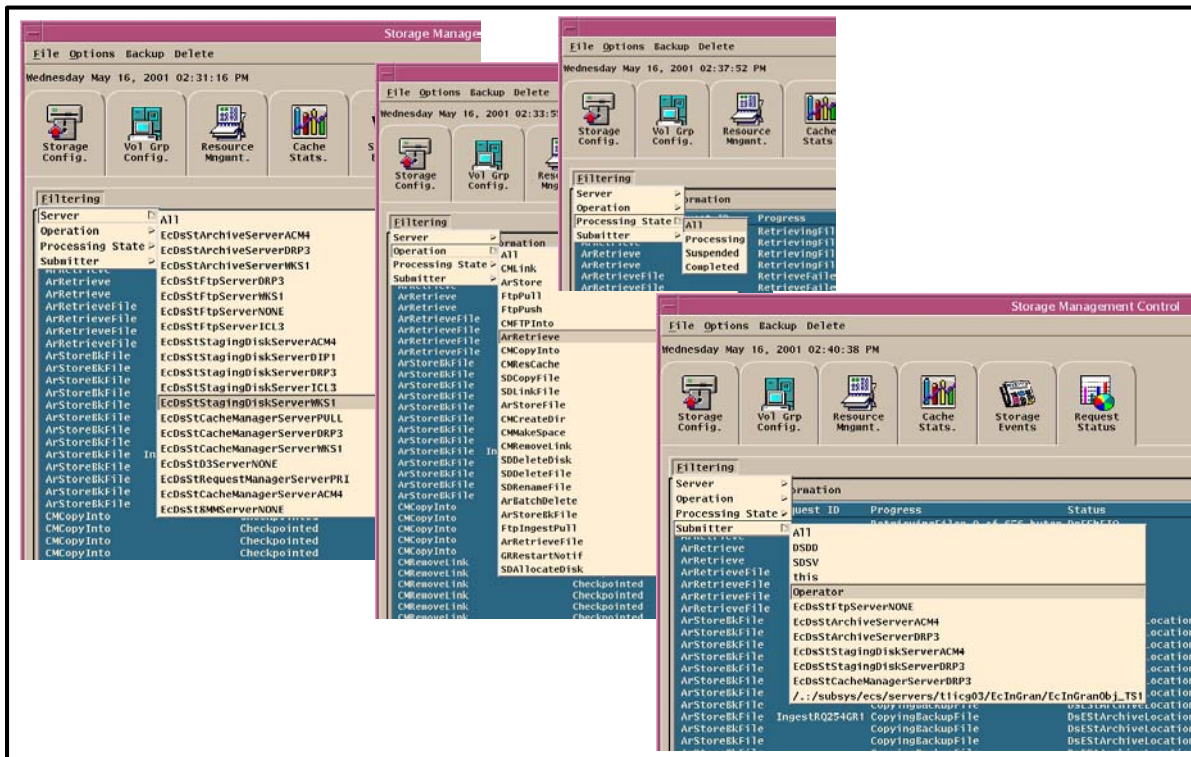


Figure 17. Storage Management Request Filtering Menus

Using the Storage Management Request Status Window to Monitor Archive Requests

- 1 Launch the DSS Storage Management GUI using UNIX commands (see procedure **Launching DSS GUIs Using UNIX Commands** [previous section of this lesson]).
 - The DSS Storage Management GUI is displayed.
- 2 Click on the **Request Status** tab.
 - The **Request Status** tab is displayed.
- 3 Observe information displayed on the **Request Status** tab of the **Storage Management Control** GUI.
 - The **Request Status Information** table displays the following information:
 - Operation.
 - Request ID.
 - Progress.
 - Status.
 - Priority.
 - When Submitted.
 - Last Updated.

- By default all storage management server requests for the last 24 hours are shown in the **Request Status Information** table of the **Request Status** tab.
 - Clicking on any of the column headers of the **Request Status Information** table causes the listed requests to be sorted in order by the column selected.
 - For example, clicking on the **Last Updated** column header causes requests to be listed in order from the least recently updated to the most recently updated.
 - The **Operator Messages** field at the bottom of the GUI displays messages concerning events occurring in storage management operations.
 - Note that storage management servers control virtually all data inserted into or retrieved from the archive; the resulting large amount of activity on the **Request Status** tab may make it useful to restrict the number of requests displayed by applying a filter (see next step).
- 4 To filter the list of requests, use the **Filtering** pull-down menu above the top left corner of the **Request Status Information** table, selecting as desired to display requests associated with a particular **Server**, **Operation**, **Processing State**, or **Submitter**.
 - The list of requests displayed in the **Request Status Information** table is restricted by the filtering choice.
 - 5 Observe the Storage Management requests displayed in the **Request Status Information** table.
 - The **Progress** and **Status** column entries in the table may provide indication for particular requests of potential problems or conditions requiring attention.
 - 6 Repeat Steps 4 and 5 as necessary to monitor Storage Management requests.
 - 7 To **exit**, follow menu path **File→Exit**.
-

Distrib'n Requests Window

The Distrib'n Requests window of the Data Distribution Operator GUI, illustrated in Figure 18, displays detailed information on individual data distribution requests and provides the capability to filter requests, change priority of requests, and designate shipping status. The window contains a list of data distribution requests that can be sorted by column. To change the priority of a selected request, select the desired priority and click on the **Apply** button in the **Change Priority** area. A selected request can be marked to indicate that it has been shipped. An **Abort** button is used to cancel a selected request. Distribution requests can also be filtered by attributes, using the **Distribution Filter Requests** window shown in Figure 19.

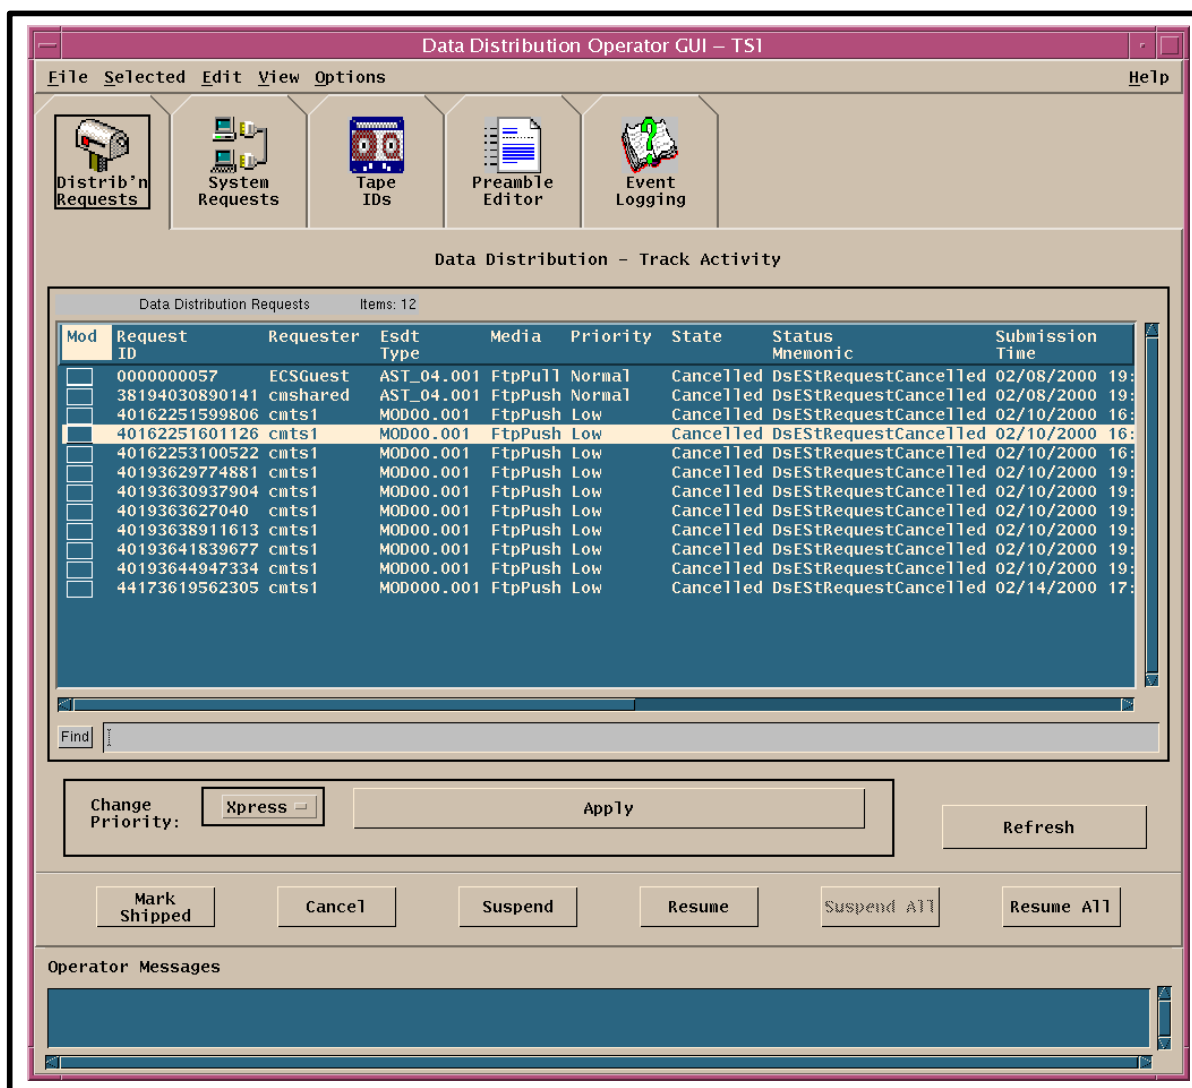


Figure 18. DDIST Distrib'n Requests Window

Distribution Filter Requests

☐ Request ID
☐ Requester
☒ All Requests

Media Type:

8MM
CDROM
D3
DLT
FtpPull
FtpPush

State:

<input checked="" type="checkbox"/> Pending	<input checked="" type="checkbox"/> Suspended
<input checked="" type="checkbox"/> Active	<input checked="" type="checkbox"/> Suspended with Errors
<input checked="" type="checkbox"/> Staging	<input checked="" type="checkbox"/> Waiting for Shipment
<input checked="" type="checkbox"/> Transferring	<input checked="" type="checkbox"/> Shipped
<input checked="" type="checkbox"/> Cancelled	<input checked="" type="checkbox"/> Failed

Figure 19. Distribution Filter Requests Window

The next section addresses retrieval of data from the archive and how to monitor distribution requests using the **Data Distribution GUI** tool and the associated **Filters**.

Retrieve Data From The Archive

Retrieval of data from the archive is a largely automated process that occurs in response to Data Distribution requests. There are a number of possible sources for Data Distribution requests:

- data orders from scientists or other ECS end users.
 - one-time orders.
 - standing orders placed as subscriptions for acquiring data.
- data requests from other ECS sites.
 - cross-DAAC orders for end users.
 - data needed as input for processing at other sites (subscriptions placed for ingest by those sites).
- internal requests for data needed for processing.

Monitoring Distribution Requests

Placing orders and subscriptions on behalf of scientists or other ECS end users is typically done by User Services representatives. Procedures for these activities are addressed in the training materials for User Services. As ECS responds to these requests, and to requests from other ECS sites or internal processes, the Archive Manager or other operators can monitor the progress of the distribution requests. Suppose that a User Services representative at your DAAC asks you to check on the status of a data distribution request from a user named Ivan Ohrdurr. The following procedure is applicable.

Monitor Distribution Requests Using the Data Distribution GUI

- 1 Launch the **DDIST GUI** using UNIX commands (see procedure **Launching DSS GUIs Using UNIX Commands** [previous section of this lesson]).
 - The **Data Distribution GUI** tool is displayed.
- 2 Click on the **Distrib'n Requests** tab.
 - The Distribution Requests window is opened.
 - A list of requests is displayed.
- 3 Click on the **Filter** push button.
 - The **Distribution Filter Requests** window opens.
 - Three filter types are displayed: **Request ID**, **Requester**, and **All**.

- 4 Click on the **Requester** button, in the radio box.
 - The cursor moves to the selection field to the right of the **Requester**.
 - 5 Enter the requester's name (in this case, **Ohrdurr**) in the text entry field opposite the **Requester** button and label.
 - 6 Click on the **All** button in the **Media Type:** area.
 - All of the Media Type toggle buttons show as selected (depressed).
 - If you are seeking only requests for a particular set of media, you can select just the button(s) for that set instead of clicking on the **All** button.
 - 7 Click on the **All** button in the **State:** area.
 - All of the State toggle buttons show as selected (depressed).
 - If you are seeking only requests in a particular state or states, you can select just the button(s) for the desired state(s) instead of clicking on the **All** button.
 - 8 Click on the **OK** push button, located at the bottom of the window.
 - The other push buttons located at the bottom of the window are **Apply**, **Cancel**, and **Help**.
 - The **Filter Requests** window is closed.
 - The Distribution Requests screen shows any requests that meet the filter criteria in the **Data Distribution Requests** field.
 - 9 If necessary, use the scroll bar at the bottom of the **Data Distribution Requests** field to scroll horizontally to view the state of the Ohrdurr request(s).
-

Deleting Granules

The system provides a **Granule Deletion** capability, complementing the automatic, scheduled deletion capability that permits operators to delete products produced and archived by the Planning and Data Processing subsystems on a scheduled basis (e.g., deletion at a certain time (configurable by the operator) after product creation.

The **Granule Deletion** capability allows operators to delete products on demand. There are a variety of circumstances that may require deletion on demand, such as:

- New PGE versions have been created and are used to reprocess large amounts of past data, creating new ESDT versions. As reprocessing progresses, operations deletes the granules for the old ESDT versions from the archive and inventory.
- It is determined that certain lower-level (e.g., Level 2) products are of little or no interest to the science or public user community. In concert with the science teams, DAAC operations personnel decide to remove these products from the inventory. Since the products are still referenced by higher-level products as inputs, the DAAC decides to keep the inventory records for production history purposes.
- One or more granules were found defective and were reprocessed on an individual basis. When the reprocessing is complete, the operator wishes to delete the old, defective granule(s) from the inventory.
- A DAAC has extended ECS with subsetting services. The subsetting products are produced outside ECS, but are then inserted into the ECS archive to take advantage of the ECS distribution capability. The DAAC writes a script to delete the subsetting products on a regular basis.

Deletion Capability and Features

The Science Data Server has provided an application programming interface (API) for deleting granules from the archive, or from both the archive and inventory since earlier releases, but the Granule Deletion capability adds a front-end command-line utility that provides several ways for selecting granules for deletion. Confirmation is generally required so that granules are not inadvertently deleted. However, the confirmation may be suppressed so that operators can run regularly scheduled deletion scripts using background execution. This suppression possibility presents an opportunity for inadvertent loss of data and so must be used with care and only after thorough testing of any deletion script.

The Science Data Server captures deletions and related errors in the application log. Operators may also specify a separate and independent delete log for immediate analysis of the success or failure of a delete operation.

Deletion Sequence

The deletion of granules from the archive involves three elements, and therefore actually occurs in stages. Two of the elements are parts of the Science Data Server (SDSRV), and the third is a part of the Storage Management (STMGT) software and Graphical User Interface (GUI).

- *Logical Deletion:* For the first stage, a command-line delete utility specifies selection criteria for deletion of granules and "logically" deletes from the inventory those granules that meet the criteria. These granules are flagged as 'deleted' and can no longer be accessed, but their inventory entries are not yet removed. The logical 'deletion' may specify, via command line input, removal of granule files from the archive (*Delete From Archive*, or DFA) only, leaving the inventory record, or it may specify *Physical Deletion*, which entails removal of the inventory record as well as removal of the files from the archive. The deletion flag consists of records in the SDSRV database. Specifically, in the DsMdGranules table, the value of the DeleteFromArchive entry is changed from N to Y, and the granule is entered in the DeletedGranules table with a time stamp recording the logical deletion time.
- *Physical Deletion:* The second stage is actual deletion from the inventory of those granules marked for physical deletion (not DFA only), which occurs when the operations staff runs the physical deletion cleanup utility script. For Physical Deletion, the script removes all inventory rows for granules that were flagged as 'deleted,' including rows referencing related information (e.g., QA data). The script writes to the STMGT database (and therefore must be run under a log in by *sdsrv_role* with authorization to write to that database), creating entries in the DsSdPendingDelete table for granules to be deleted. This includes entries for granules that are to be physically deleted, as well as those designated DFA only. The operations staff controls the lag time between logical deletion and physical deletion. That lag time is entered into the physical deletion script, which deletes only inventory entries for granules that have been logically deleted prior to that time period.
- *Deletion from Archive (DFA):* STMGT provides a GUI screen, as illustrated in Figure 20, that allows the operator to initiate the removal from the archive of the files listed in its deletion table (populated by SDSRV). STMGT creates requests to the archive servers to delete files. The STMGT GUI can be used to look at the state of the deletion requests. Files that are successfully deleted have their associated rows removed from the STMGT database table.

Periodically, as sufficient data removal from the archive makes it appropriate, operations may elect to reclaim the tape space and recycle archive tapes. The AMASS software commands (*volcomp*, *volclean*, *volformat*, *volstat*) are used for that purpose.

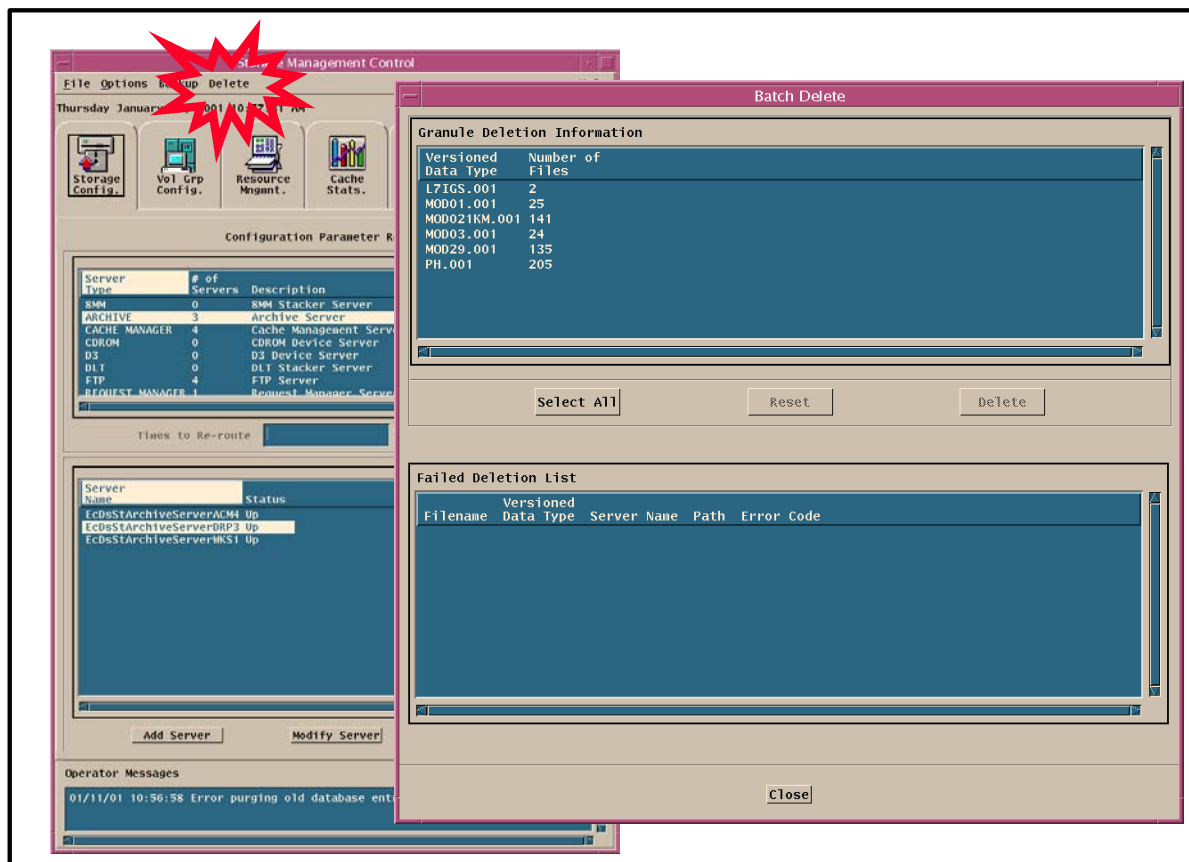


Figure 20. STMG GUI screen for granule deletion from archive

Suppose you learn from one of the production monitors at your DAAC that a defective granule inserted during Data Processing on January 12, 2001 at 11:03:42 (Granule ID is :SC:MOD29.001:5936:1.HDF-EOS) has been replaced through reprocessing. Use the following procedure for deleting the defective granule from the inventory and archive.

Delete Granule from the Inventory and Archive

1 Access the command shell.

- The command shell prompt is displayed.

NOTE: Commands in Steps 2 through 11 are typed at a UNIX system prompt.

2 Type **setenv DISPLAY *clientname*:0.0** and then press the **Return/Enter** key.

- Use either the terminal/workstation IP address or the machine-name for the *clientname*.

- 3 Start the log-in to the SDSRV client server by typing `/tools/bin/ssh hostname` (e.g., `e0acs05`, `g0acs03`, `n0acs04`, or `l0acs04`) and then press the **Return/Enter** key.
 - If you receive the message, **Host key not found from the list of known hosts. Are you sure you want to continue connecting (yes/no)?** type yes (“y” alone does not work).
 - If you have previously set up a secure shell passphrase and executed `sshremote`, a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears; continue with Step 4.
 - If you have not previously set up a secure shell passphrase; go to Step 5.
- 4 If a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears, type your *Passphrase* and then press the **Return/Enter** key. Go to Step 6.
- 5 At the `<user@remotehost>'s password:` prompt, type your *Password* and then press the **Return/Enter** key.
- 6 To run the Granule Deletion Client specifying ESDT ShortName, ESDT version, and granule insert time coverage, type the command
`/usr/ecs/<MODE>/CUSTOM/utilities/EcDsGranuleDeleteClientStart -name MOD29 -version 001 -insertbegin 01/12/2001 11:03:00 -insertend 01/12/2001 11:04:00 -log /usr/ecs/<MODE>/CUSTOM/logs/GranDel1.log -physical ConfigFile /usr/ecs/<MODE>/CUSTOM/cfg/EcDsGranuleDelete.CFG ecs_mode <MODE>`, and then press the **Return/Enter** key.
 - The client executes and displays the number of granules for deletion (in this case, 1), and prompts the user **Do you want to continue [y/n]?**.
 - **NOTE:** The Granule Delete tool provides other ways to delete granules. The deletion specified here could be achieved using one of these ways. For example, if the granule time coverage is available, the command can specify that instead of the insert time. The tool also permits referencing a list of granules by geoid in a file created for the purpose of providing that list as input to the tool. The desired deletion could be achieved by creating a file (e.g., `dbids1.in`) containing the geoid(s) for the granule(s) to be deleted (in this case, `SC:MOD29.001:5936`) and typing the command `EcDsGranuleDeleteClientStart -geoidfile dbids1.in -log /usr/ecs/<MODE>/CUSTOM/logs/GranDel1.log -physical ConfigFile /usr/ecs/<MODE>/CUSTOM/cfg/EcDsGranuleDelete.CFG ecs_mode <MODE>`. The file can contain geoids for multiple granules. Similarly, the deletion could be achieved by creating a file containing the local granule Id and using a command specifying `... -localgranulefile locgrn1.in ...`.

- 7 Type **y** and then press the **Return/Enter** key.
 - The process continues to completion.
 - **Note:** The deletion actions are displayed in the Deletion log and in the Science Data Server ALOG, including information on the user ID of the requester, the ShortName, VersionID, and granule insert time of the request. In addition, the EcDsGranuleDelete.ALOG may contain useful information. It is also possible to view the SDSRV database to verify the granule tagging for deletion; the granule should appear in the database with a value of the **time the tool was run** for DeleteEffectiveDate, and a value of **N** for DeleteFromArchive. The DsMdDeletedGranules table should list the granule just requested for deletion.
- 8 To run the Deletion Cleanup Utility, type **EcDsDeletionCleanup.pl** and then press the **Return/Enter** key.
 - The script prompts **Enter lag time in days:**.
- 9 Type **0** and then press the **Return/Enter** key.
 - The script prompts **Is this correct? [y/n]**.
 - **NOTE:** In this training exercise, this step specifies a lag time of **0** and the next step confirms the entry as correct. This lag time is used in the exercise to illustrate the functioning of the tool within the time allotted for training. Typically, a lag time of **0** is not used; instead, it is recommended to use a lag time of **30** days, allowing time when the data are not accessible but before the data are physically removed to identify possible deletion errors.
- 10 Type **y** and then press the **Return/Enter** key.
 - The script prompts **Enter mode of operation:**.
- 11 Type **<MODE>**, where **<MODE>** is the mode in which you are making the deletion (typically **OPS**, **TS1**, or **TS2**) and then press the **Return/Enter** key.
 - The script prompts **Enter log file name:**.
- 12 Type **DelCleanup1.log** and then press the **Return/Enter** key.
 - The script prompts **Enter Sybase User:**.
 - **NOTE:** It is possible to press the **Return/Enter** key without typing a name for the log file; in response, the script returns **Using default Log File name (Default is DeletionCleanup.LOG)**. before prompting **Enter Sybase User:**.
- 13 Type **sdsrv_role**, and then press the **Return/Enter** key.
 - The script prompts **Enter Sybase User Password:**.

- 14 Type **<password>**, where **<password>** is the Sybase password (*Note*: This step may require action by the Database Administrator).
 - The script prompts **Enter Sybase SQL Server Name:**.
- 15 Type **<x>0acg<nn>_srvr** (**e0acg11_srvr** at EDC, **g0acg01_srvr** at GSFC, **l0acg02_srvr** at LaRC, or **n0acg01_srvr** at NSIDC) and then press the **Return/Enter** key.
 - The script prompts **Enter SDSRV's database name:**.
- 16 Type **EcDsScienceDataServer1_<MODE>**, where **<MODE>** is the mode in which you are making the deletion (typically **OPS**, **TS1**, or **TS2**) and then press the **Return/Enter** key.
 - The script prompts **Enter STMGT's database name:**.
- 17 Type **stmgtdb1_<MODE>**, where **<MODE>** is the mode in which you are making the deletion (typically **OPS**, **TS1**, or **TS2**) and then press the **Return/Enter** key.
 - The script prompts **Enter Batch Size (10000):**.
- 18 Type **1** and then press the **Return/Enter** key.
 - The Deletion Cleanup Utility script displays a list of actions as they are completed, and then displays the number of granules to be deleted from the archive (DFA) and physically deleted, with a confirmation prompt **Do you wish to continue deleting these granules? [y/n]**. All granules in the DeletedGranules table are displayed because the entered lag time of 0 specifies deletion of all granules tagged for deletion.
- 19 Type **y** and then press the **Return/Enter** key.
 - Execution of the Deletion Cleanup Utility script completes.
 - *Note*: In the SDSRV database, the SDSRV Staging table (DsMdStagingTable) can be observed for transfer of data to the STMGT database (in increments of the specified batch size, in this case **1**); when the transfer is complete, the table is empty. In the STMGT database, the STMGT Pending Delete table (DsStPendingDelete) can be observed for receipt of the data; all granules specified in the delete request are received. The Deletion Cleanup log displays messages about the actions, indicating that information is placed in the STMGT database in increments of the specified batch size, in this case **1**.
- 20 On the STMGT GUI, to view the ESDTs with granules targeted for deletion, follow menu path **Delete→Batch Delete**.
 - The **Batch Delete** window is displayed, listing the number of files for each ESDT/Version pair tagged for deletion in the **Granule Deletion Information** field.
- 21 To select data for deletion from the archive, click on an ESDT/Version pair (in this case, **MOD29.001**).
 - The selected ESDT/Version pair is highlighted.

- 22 Click the **Delete** button.
- A confirmation Delete Warning prompt asks **Are you sure you want to delete the selected files?**
- 23 To confirm the deletion, click the **OK** button.
- The delete request continues to completion. If you close the **Batch Delete** window (by clicking on the **Close** button) and then re-launch the window (by following menu path **Delete→Batch Delete**), the ESDT/Version pair no longer appears in the **Granule Deletion Information** field of the **Batch Delete** window.
 - *Note:* The delete actions can be tracked via messages in the Archive Server log files (EcDsStArchiveServer.ALOG, EcDsStArchiveServerDebug.log)
-

Undelete Capability

In the event that it is desirable to restore granules that have been marked for deletion (although not yet actually removed physically from the archive), the Granule Deletion Tool provides an undelete capability. This is implemented by command line options for use with the command **EcDsGranuleDeleteClientStart** and its selected options as follows:

- **-undelphysical** to undelete granules that have been marked logically deleted by the **-physical** option;
- **-undelDFA** to undelete granules that have been marked logically deleted by the **-DFA** option;
- **-displayUndelPhysical** to display granules that will be undeleted with the **-undelphysical** option and other selected options (*Note:* the number of granules returned with this option is limited by MaxGeoidFileLines and MaxCollectorSize parameters in the EcDsGranuleDelete.CFG file and/or the DBMAXRESULTS parameter for the Science Data Server);
- **-displayUndelDFA** to display granules that will be undeleted with the **-undelDFA** option and other selected options (*Note:* the number of granules returned with this option is limited by MaxGeoidFileLines and MaxCollectorSize parameters in the EcDsGranuleDelete.CFG file and/or the DBMAXRESULTS parameter for the Science Data Server).

These options are used in conjunction with other command line options of the Granule Delete tool. For example, it is possible to undelete or display granules using the command line options for selection by ESDT short name, version, and granule time coverage, or to undelete or display granules using the command line option to specify a separate input file.

Loading Archive Media

Automatically Loading Archive Media

For the STK storage facility, each Powderhorn is equipped with a 21-tape Cartridge Access Port (CAP). In automatic mode, tapes may be placed in the CAP for automatic loading. Tapes are also ejected through the CAP when identified for ejection using a command at the host for the STK Automated Cartridge System Library Software (ACSL).

The following procedure addresses automatic media loading.

Automatically Loading STK Powderhorn Archive Media

- 1 Log in as **amass** or **root** at the FSMS SGI host (workstation **x0drg##**).
 - NOTE: The **x** in the workstation name will be a letter designating your site: **g** = GSFC, **m** = SMC, **l** = LaRC, **e** = EDC, **n** = NSIDC, **o** = ORNL, **a** = ASF, **j** = JPL; the **##** will be an identifying two-digit number (e.g., **n0drg01** indicates an FSMS SGI server at NSIDC).
- 2 At the FSMS host, type **/usr/amass/bin/bulkinlet SP** and then press the **Return/Enter** key.
 - The Cartridge Access Port (CAP) door unlocks (audible unlatching sound).
 - *Note:* If you have removed an existing volume and are re-inserting it, do not use the **SP** option, which puts the volume in the general space pool. Instead type **/usr/amass/bin/bulkinlet <volgrp>**, where **<volgrp>** is the volume group from which the volume was removed. This will put the volume back where it was before removal.
- 3 Write down or note the bar code number(s) on the label(s) of the cartridge(s), open the recessed latch on the CAP door and insert the tape(s), solid black side up, with the bar code label facing you, and close the door.
 - The robot scans all the volumes.
 - Data for the newly inserted media are displayed, including bar codes, associated volume numbers, and, in the **flag** column, the letters **IUO**, indicating that the volumes are inactive (**I**), unformatted (**U**), and offline (**O**).
- 4 For any newly inserted media, it is necessary to issue a formatting command. For the new 9940 tapes, type **/usr/amass/bin/volformat -b 256k ###**, where **###** is the volume number, and then press the **Return/Enter** key. You can enter more than one, separating each number from the preceding one with a space.
 - A message requests confirmation that you wish to continue.

- 5 Type **y** and then press the **Return/Enter** key.
 - A message is displayed requesting further confirmation, stating that **The following volumes will be formatted:** and listing volume numbers, followed by **(Y-N)**.
 - 6 Type **y** and then press the **Return/Enter** key.
 - After a few minutes, a message **Completed formatting all volumes** is displayed.
 - 7 To verify that the volume(s) are inserted, type **/usr/amass/bin/vollist** and then press the **Return/Enter** key.
 - Data for the media are displayed; the **flag** column shows that the newly formatted volumes are inactive (**I**).
 - 8 To activate the media for use, type **/usr/amass/bin/volstat** and then press the **Return/Enter** key.
 - Data for the media are displayed; the **flag** column shows that the volumes are now active (**A**).
-

Manually Loading Archive Media

With the **bulkload** command, you bypass the CAP and manually load media directly into the library bins. Typically, this will only be done at the initial load of the system with large numbers of media volumes. The **bulkload** command enables AMASS to determine what type of media have been placed in the library and to convey this information to the AMASS database. The following procedures are applicable.

Manually Loading STK Powderhorn Archive Media Procedure

- 1 To manually insert a tape into the Powderhorn, login to the control software (ACSL) using the **acssa** account.
- 2 Type **enter 0,0,0** and then press the **Return/Enter** key.
 - The Cartridge Access Port (CAP) door unlocks (audible unlatching sound).
- 3 Write down or note the bar code number(s) on the label(s) of the cartridge(s), open the recessed latch on the Cartridge Access Port (CAP) door and insert the tape(s), solid black side up, with the bar code label facing you, and close the door.
 - The robot scans all the volumes.
- 4 At the AMASS host, type **bulkload -s SP** and then press the **Return/Enter** key.
 - The AMASS database is populated with data for the volumes in the library, including bar codes, associated volume numbers, and status -- inactive (**I**), unformatted (**U**), and offline (**O**). The data may be reviewed using the **vollist** command.

- *Note:* If you are loading a very large number of volumes, such as at initial load, and choose to bypass the CAP and place the volumes directly in the LSM slots, data about the volumes will not be immediately available to ACSLS for communication to AMASS. You will first have to use the ACSLS **audit** command to initiate an audit of the LSM, a process that may take several hours.

Caution

Inactivate AMASS before using the following command.

- 5 To view a list of media in the library, type **medialist -3**, and then press the **Return/Enter** key.
 - The **-3** option indicates the STK Powderhorn.
 - The utility reads the library element status stored in the library, and information about the library contents, including the status (**FULL** or **EMPTY**) of the elements.
-

Formatting a Volume

Volume formatting was part of the procedure for automatically loading media (previous section of this lesson). However, it may be necessary or desirable at some time to format a volume that is already present in the archive. To format a volume, it must be online. A volume is placed online using the **volloc** command. Formatting a volume destroys any files on that volume. Before formatting a volume, check to make sure it does not have any files that should be saved. The following procedure is applicable.

Formatting a Tape Volume

- 1 To put the volume online, at the FSMS host, type **/usr/amass/bin/volloc -n ###**, and then press the **Return/Enter** key.
 - **###** is the number of the volume.
- 2 To verify there are no files on volume, type **/usr/amass/bin/volfilelist <Vol. No.>**, and then press the **Return/Enter** key.
 - No files are displayed.
 - If a list of files is returned, indicating that the volume is not empty, before proceeding verify that you have the correct volume and that it is to be formatted.
- 3 To format the volume, type **/usr/amass/bin/volformat -b 256k ###**, and then press the **Return/Enter** key.
 - **###** is the number of the volume.

- 4 To verify status of the volume, type **/usr/amass/bin/volprint -a ###**, and then press the **Return/Enter** key.
 - ### is the number of the volume.
-

Remove Media

To remove media from the archive, use the following procedures.

Remove Media from STK Powderhorn

- 1 Determine which volumes you want to remove by utilizing the volume number. If necessary to review volume numbers and other information, log into the AMASS host, type **/usr/amass/bin/vollist** and then press the **Return/Enter** key.
 - 2 If there are only a few volumes to remove, from AMASS, for each volume to be removed type **/usr/amass/bin/voloutlet ###**, where ### is the volume number, and then press the **Return/Enter** key.
 - AMASS marks the volume off-line and the volume is transferred to the CAP.
 - 3 For the STK Powderhorn, open the recessed latch on the Cartridge Access Port (CAP) door and remove the tape(s)
-

Backup Archived Data

The ECS archive design incorporates programmed backups of archived data. System requirements specify that a percentage of archived data be duplicated for local and off-site storage to provide for data safety. However, the large volume of ECS archived data merits finding alternatives to complete backup of all volumes in the libraries. Selection of data for backup is based on assessment of the feasibility of recovery in the event of data loss.

It is imperative to backup data that would be irretrievable if lost. Such data are saved to the archive, saved to local backup, and saved to offsite backup. Many data elements that will be archived, however, could be retrieved in the event of loss. For example, in the event of loss of a higher level product that is an output of processing a lower level product, it would be possible to restore the higher level product by reprocessing the lower level product. As another example, ECS will often archive a lower level product from a data provider, but that product may also be retained in the archives of the data provider. If the product were lost from the ECS archive, it would be possible to ingest it again from the data provider, using appropriate Ingest procedures.

Thus, when data are inserted into the archive (e.g., through Ingest, from Processing), up to three copies of the data may be created, reflecting different types of data use:

- the active archive copy, available for distribution or other use (volume group is specified in the *Archive ID*).
- a copy to be retained for local backup (volume group is specified in the *Backup ID*, typically created by appending "B" to the ShortName and VersionID; e.g., AST_L1B.001B).
- a copy to be sent to offsite backup storage (volume group is specified in the *Offsite ID*, typically created by appending "O" to the ShortName and VersionID).

Creating Offsite Backups

The paths for creation of the data copies are specified for each ESDT when it is loaded, using the **Add Volume Group** dialog accessible from the **Vol.Grp. Config.** tab on the Storage Management GUI as illustrated in Figure 21. The Archive ID (for the archive copy) and the Backup ID (for the local backup copy) should reflect different archives if possible (i.e., different LSMs), to spread the risk of loss. The Offsite ID will not be a remote site path, but rather a local path for making copies to be sent for offsite storage. The requirements to implement creation of offsite backups (see Document 611-CD-610-002, *Mission Operation Procedures for the ECS Project*, section 17.6.1.1) include:

- creating a subdirectory and volume group for offsite backups.
- using the **Vol.Grp. Config.** tab of the Storage Management GUIs to add the volume group to the appropriate archive server and set the offsite ID to be the three-character specification for the local site (e.g., EDC, GSF, LAR, NSC).
- updating the STMGT database offsite table with the volume group name created for offsite storage (may require assistance of Database Administrator).
- adding volumes to the volume group as needed.

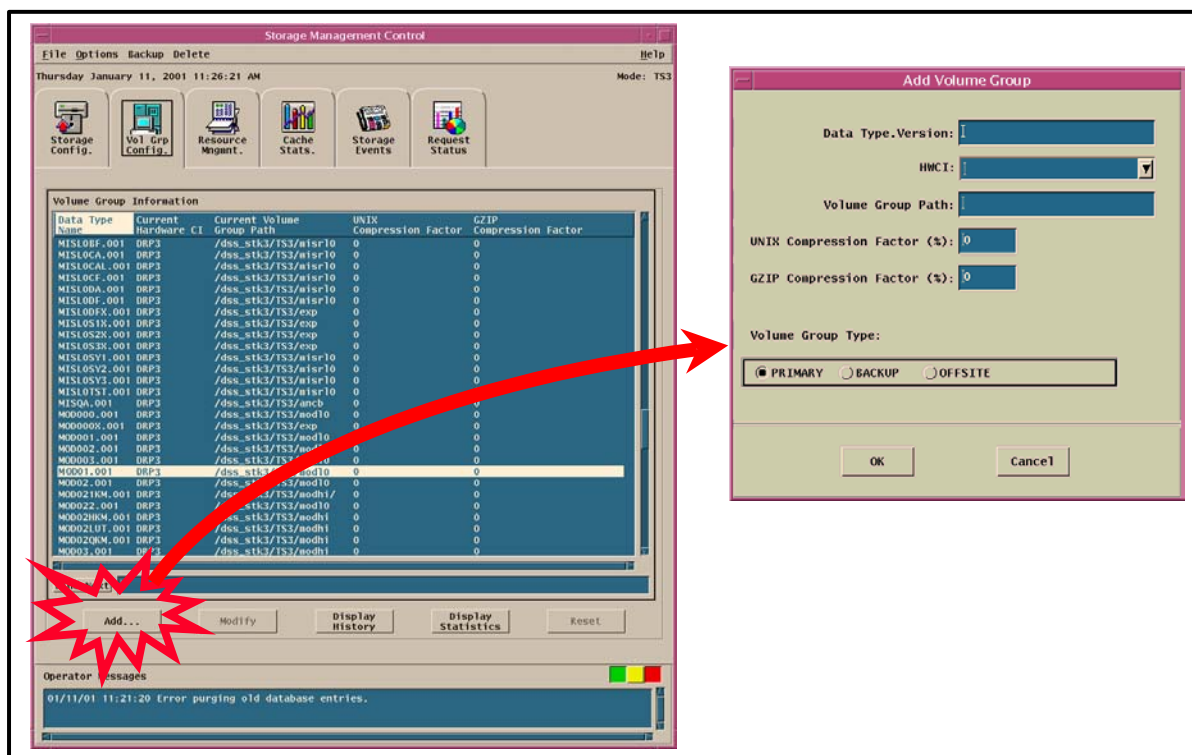


Figure 21. Storage Management, Add Volume Group dialog

Each site is responsible for arranging its own secure offsite storage. The offsite backup cartridges are removed from the archive storage facility using procedures already described. For local and/or offsite storage of specific archive data, the DAAC Archive Manager generates or directs the generation of a list of selected data. At the time the files are archived, they are written to specific volume groups that correspond to the three data usage types identified in the preceding paragraph. Only files belonging to the data usage type are written to the tapes in a specific volume group. Hence, the Archive Manager can determine the tapes that should be stored for local backup and those for offsite storage. This can be accomplished using the AMASS administration **vollist** command.

vollist - this command lists all volumes and their current status information. It will optionally accept a specific volume number. Information displayed includes: volume number, volume group, jukebox number, slot position, volume label/bar code, current volume status, amount of used space on the volume, amount of available space on volume, percentage of space no longer referenced, and number of read or write errors on volume.

If there are other files designated for local and offsite backup which have not been written to a specific volume group, the Archive Manager can use the appropriate AMASS administration commands, **dirfilelist** and/or **volfilelist**, to locate the appropriate archive volume that contains the designated archived files.

dirfilelist: this command lists the files under a directory. One directory or file is displayed on a line. Information displayed includes: file name, volume number, starting block number, file size, file permissions, number of hard links, numeric user id, numeric group id, last time file was accessed, and last time file was modified.

volfilelist: this command lists all of the files on a volume and accepts as input a volume number. One file is displayed on a line. Information displayed includes: file name, volume number, starting block number, file size, file permissions, number of hard links, numeric user id, numeric group id, last time file was accessed, and last time file was modified.

For more information about these AMASS commands, refer to the *AMASS System Administrator's Guide*. The following procedure is applicable for creating offsite backups for science data.

Creating Offsite Backups for Science Data

- 1 Launch the DSS Storage Management GUI using UNIX commands (see procedure **Launching DSS GUIs Using UNIX Commands** [previous section of this lesson]).
 - The DSS Storage Management GUI is displayed.
- 2 Click on the **Vol Grp Config** tab to display the Volume Group information.
 - The **Vol Grp Config** tab information is displayed.

- 3 Click on the **Add . . .** button below the Volume Group Information field.
 - The **Add Volume Group** window is displayed.
- 4 In the **Add Volume Group** window, click in the **Data Type.Version:** field.
 - The cursor moves to the **Data Type.Version:** field.
- 5 Type the ESDT *ShortName* and *Version* (e.g., MOD01.001) of the data type for which the volume group is to be created.
 - The typed entry appears in the **Data Type.Version:** field.
- 6 In the **Add Volume Group** window, click on the pull-down arrow at the end of the **HWCI:** field.
 - A pull-down menu displays designators of the hardware configuration items available for storing data.
- 7 Click on the designator for the hardware configuration item where the archive copies of data for the ESDT are to be stored.
 - The selected designator is displayed in the **HWCI:** field.
- 8 In the **Add Volume Group** window, click in the **Volume Group Path:** field.
 - The cursor moves to the **Volume Group Path:** field.
- 9 Type the full path identification for the storage of active archive data for the ESDT (typically, the path will be of the form **dss_stk***n*/**<MODE>**/**xxxxx**, where *n* is a number designating a StorageTek Library Storage Module, **MODE** is **OPS**, **TS1**, or **TS2**, and **xxxxx** is a short identifier for what is being stored; e.g., **dss_stk1/OPS/modl0**).
 - The typed entry appears in the **Volume Group Path:** field.
- 10 In the **Volume Group Type:** radio box, click on the **PRIMARY** button.
 - The button depressed appearance indicates selection of **PRIMARY**, signifying that the volume group being created is for primary storage for active archive use.
- 11 Click on the **Save and Add Next VG** button at the bottom of the **Add Volume Group** window.
 - The volume group is created for display in the **Volume Group Information** field on the **Vol Grp Config.** tab of the Storage Management GUI.
- 12 In the **Add Volume Group** window, click in the **Volume Group Path:** field.
 - The cursor moves to the **Volume Group Path:** field.

- 13 Change the data entered at Step 9 to identify the full path for the storage of local backup data for the ESDT.
 - *Note:* This step is only for those ESDTs that require local backup.
 - The typed entry appears in the **Volume Group Path:** field.
- 14 In the **Volume Group Type:** radio box, click on the **BACKUP** button.
 - The button depressed appearance indicates selection of **BACKUP**, signifying that the volume group being created is for storage for local backup use.
- 15 Click on the **Save and Add Next VG** button at the bottom of the **Add Volume Group** window.
 - The volume group is created for display in the **Volume Group Information** field on the **Vol Grp Config.** tab of the Storage Management GUI.
- 16 In the **Add Volume Group** window, click in the **Volume Group Path:** field.
 - The cursor moves to the **Volume Group Path:** field.
- 17 Change the data entered at Step 13 to identify the full path for the creation and initial storage of offsite backup data for the ESDT.
 - *Note:* This step is only for those ESDTs that require offsite backup.
 - The typed entry appears in the **Volume Group Path:** field.
- 18 In the **Volume Group Type:** radio box, click on the **OFFSITE** button.
 - The button depressed appearance indicates selection of **OFFSITE**, signifying that the volume group being created is for creation and initial storage for offsite backup use.
- 19 Click on the **Save and Exit** button at the bottom of the **Add Volume Group** window.
 - The volume group is created for display in the **Volume Group Information** field on the **Vol Grp Config.** tab of the Storage Management GUI.
 - Data stored in the volume group for OFFSITE backup can be safeguarded by removing tapes that have data stored on them (see procedure 17.2.4) and transporting the tapes to a secure offsite storage location.

Creating a Backup for AMASS

The Archive Manager should periodically create a backup to guard against loss of the AMASS database and functioning. The archive storage format used by AMASS is a proprietary format designed to optimize storage and retrieval speed. The command **vgexport -q** can be used to create a text file, storable on magnetic media, which can be used with the AMASS format archive tapes and the command **vgimport** to recover from the loss. This command exports the AMASS database for a specified volume group to standard out (**stdout**), a file containing the

directory structure and media attributes (e.g., media type, ownership, timestamp) for the volume group. The file is located in **/usr/amass/filesysdb** and is exported as standard ASCII text. Use the following procedure.

Create a Backup for AMASS

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01) as **amass** or **root**.
- 2 Type **/usr/amass/bin/vgexport -q**.
 - A file named **stdout** is created in **/usr/amass/filesysdb**.
 - *Note:* The **stdout** file is useful only with the archive volumes represented in the AMASS database.

Replacing the AMASS Database Backup Volume (Volume 1)

The AMASS database backup is stored in the archive on Volume 1. "Volume 1," hard coded to be the backup volume, actually designates one of the last volumes in the StorageTek LSM, to prevent its inadvertent use as a data volume. Whenever **amassbackup** is run, AMASS issues an e-mail message with information on volume capacity and usage. It is also possible to issue the command **vollist 1** to display how much space is left on the volume, or **volprint 1** for still more detail. If the volume becomes full *during* a backup attempt, the backup will fail and it is necessary to initialize a new backup volume and perform a full backup as described in the following procedure.

Replace a Full Backup Volume (Volume 1)

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01) as **amass** or **root**.
- 2 Type **/usr/amass/bin/voloutlet 1** and then press the **Return/Enter** key.
 - The LSM robot places the Backup Volume in the CAP.
- 3 Open the recessed latch on the CAP door; remove the Backup Volume tape and store it in a safe place.
- 4 Physically designate the new Backup Volume tape so that it can be easily discriminated from other volumes (e.g., write "Backup Volume" on the tape, color code the tape, or make and display a note of its home storage slot or preprinted barcode).
- 5 Note the pre-printed number on the volume label (e.g., 112102), insert the new Backup Volume in the CAP, and close the door.
 - The robot scans the volume.

6 At the AMASS host, type **/usr/amass/bin/bulkinlet -u** and then press the **Return/Enter** key.

- AMASS assigns the Backup Volume a unique volume number.
- AMASS marks the volume **ONLINE** in the AMASS database.
- AMASS assigns the Backup Volume to the last barcode position in the library.
- AMASS gives the volume a **BACKUP VOLUME** label.

7 Type **/usr/amass/bin/vollist 1**, and then press the **Return/Enter** key.

- AMASS displays the following:

VOL NUM	VOL GRP	JUKE NUM	POS	VOL LABEL	FLAGS	USED (MB)	AVAIL (MB)	DEAD (%)	ERRS
1	0	1		BACKUP-VOLUME	I	0	20000	0	0

8 To change the Volume Label field from **BACKUP-VOLUME** to the preprinted media number (e.g., 112102), type **/usr/amass/bin/vollabel 1 112102** and then press the **Return/Enter** key.

9 Type **/usr/amass/bin/vollist 1**, and then press the **Return/Enter** key.

- AMASS displays the following:

VOL NUM	VOL GRP	JUKE NUM	POS	VOL LABEL	FLAGS	USED (MB)	AVAIL (MB)	DEAD (%)	ERRS
1	0	1		112102	I	0	20000	0	0

10 Type **/usr/amass/bin/volformat -u** and then press the **Return/Enter** key.

- A message requests confirmation that you wish to continue.

11 Type **y** and then press the **Return/Enter** key.

- A message is displayed requesting further confirmation, stating that **The following volumes will be formatted: 1 (Y-N)**.

12 Type **y** and then press the **Return/Enter** key.

- After a few minutes, a message **Completed formatting all volumes** is displayed.

13 To verify that the volume is inserted, type **/usr/amass/bin/vollist 1** and then press the **Return/Enter** key.

- Data for the media are displayed; the **flag** column shows that the newly formatted volume is inactive (**I**).

14 Type **/usr/amass/bin/amassbackup -fv** and then press the **Return/Enter** key.

- AMASS performs a full backup with the verbose option of the AMASS database and transaction logs.

Create Replacement Backups Manually from Existing Archives

If loss of data necessitates obtaining and inserting backup data from local or off-site storage, it is necessary to create replacement data to be returned to backup storage. Use the following procedure.

Create Replacement Backups

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01) as **amass** or **root**.
 - 2 Type **/usr/amass/bin/volcopy -c <source> <destination>** (where *<destination>* is the volume number of the destination volume and *<source>* is the volume number of the source volume), and then press the **Return/Enter** key.
 - The **-c** option specifies copy of the source to the destination.
 - A bit for bit copy of the source (the cartridge to be copied) is made at the destination (an available, unused cartridge). Because the copy procedure depends on the amount of data on the source cartridge, the process can take as long as an hour to complete.
 - *Note:* After starting a **volcopy** procedure, do not attempt to kill the process with the **kill -9** command.
 - 3 A hardcopy/softcopy list of the files backed up should be created and kept for future file restoration operations.
 - 4 Remove the backup volume(s) and send to offsite storage area, as appropriate.
-

Restore Archive Data

Although Archive hardware has been selected for high reliability, there may be an occasional tape failure or drive error. In the event that such errors cause loss or corruption of the primary copy of Archive data, it will be necessary to attempt data recovery or to re-archive equivalent data. There are a number of potential means for restoring lost or corrupted data.

Use of Backup Data for Recovery

Depending upon the circumstances and nature of the loss, restoration of archive data may take the form of:

- copying from a local backup to the original or a new primary copy.
- reprocessing.
- obtaining replacement data from an external data provider.
- restoring the AMASS database and/or the ACSLS database.

In the event of catastrophic loss of the archive and local backup cartridges, it will be necessary to retrieve backups from off-site storage. A different process required to recover from a failed attempt at archive storage is addressed in a subsequent section on archive monitoring and fault notification.

Manual Data Recovery from Local or Offsite Backup Tapes

If a backup volume is available and contains the data that were lost or corrupted on the primary copy, the data can be copied using standard UNIX commands. Detailed procedures are available in section 17.6.2 of Document 611-CD-610-002, *Mission Operation Procedures for the ECS Project*. If the backup volume must be obtained from offsite storage, it must then be inserted into the archive and brought on line. The procedures for loading archive media were addressed under a preceding topic. The requirements then entail:

- using the **Storage Config.** tab of the Storage Management GUIs to view the volume groups of the appropriate archive server and to find the files in the primary and backup volume groups.
- using the UNIX copy command (**cp** or **dd**) to copy the lost or corrupted file from the backup version to the primary version.
- as appropriate (i.e., if the recovery is one of a set of files to be restored, for example, because they were lost from a damaged tape), removing the names of the files recovered from the list of files to be recovered by other means.

If an entire volume is to be copied, perform the procedure to create replacement backups as addressed under a previous topic; if recovery is from offsite, send the backup back to secure offsite storage.

Reprocessing

If it is possible to produce a lost data product by running a Product Generation Executive (PGE) on other data available in the archive, recovery of the lost product may be achieved by this reprocessing. The reprocessing will be a case of on-demand processing, for which procedures are addressed in a separate lesson on Production Planning and Processing. In this case, the resultant recovered file will have a new Universal Reference (UR) and a new Production Date and Time. As part of the input to this process, the operator needs the following information for each file to be recovered:

- the Archive unique filename for the file.
- the Archive IDs of the primary archive and backup archive.
- the file checksum if available.

Requesting Replacement Data from Provider

Where the archived data that are lost are not available in local or off-site backups, but were originally acquired from an external data provider who retains a copy of the data in the archives of the data provider, recovery may be achieved by re-ingesting the data. Most Level 0 (L0) instrument data in the ECS archive was originally supplied by the EOS Data and Operations System (EDOS). Missing or corrupt ECS L0 files can usually be restored from the EDOS archive. If ECS operations at the Goddard Space Flight Center Earth Sciences DAAC requests a specific L0 product or products, EDOS can fulfill the request by supplying DTF-2 archive tape(s) containing the EDOS Production Data Sets (PDSs). For other DAACs, without DTF-2 tape drives available, EDOS supplies replacement PDSs using an FTP interface.

Requests for replacement L0 data from EDOS are made using an ESDIS-sponsored Data Reorder web tool. The tool is accessed from an ESDIS web page listing several links for reports and tools in the End-to-End Tracking System, as illustrated in Figure 22. The specific information to be entered using the tool in a reorder may vary somewhat depending on the platform and instrument involved in the initial capture of the data. Accordingly, as illustrated in the figure, there are separate links in the **Data Reorder** row of the Related Links table on the page.

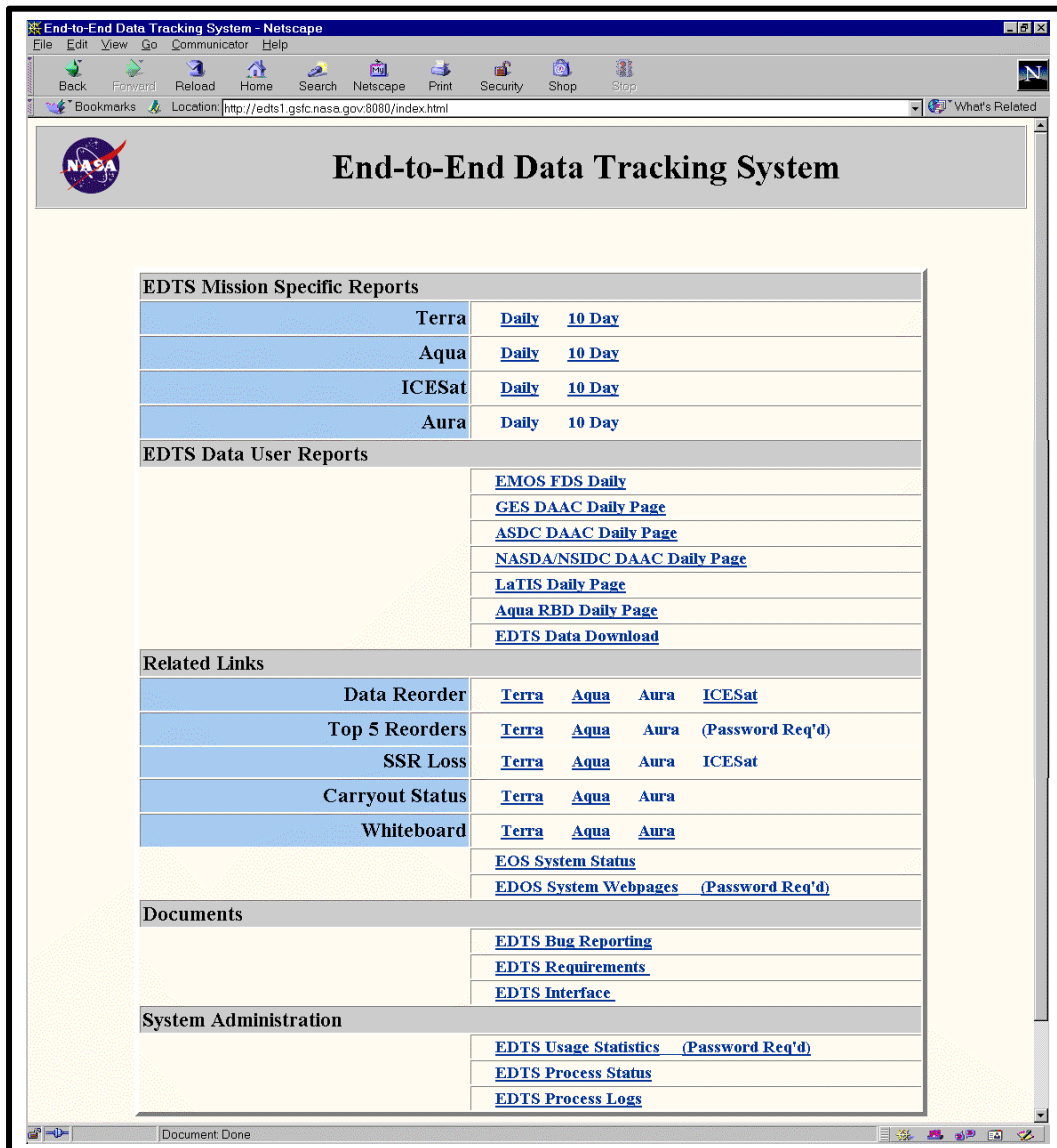


Figure 22. End-to-End Tracking System web page providing access to the Data Reorder tool

A click on the appropriate Data Reorder link results in display of a page specific to the platform specified in the link. For example, a click on the **Terra** link in the Data Reorder row of the page illustrated in Figure 22 results in display of the page shown on the left side of Figure 23. A subsequent click on the **Add New Request** link displays the page shown on the right side of Figure 23, permitting specification and submission of the reorder request.

EOS Terra Data Reorder Request

Add New Request

[Data Reorder](#) [Tapes](#)

Requestor's Name* Requestor's Organization*

Requested data (UTC)* Year: 2003 DOY: 1 Start time: 0000 Stop time: 0200

Request Reason* ☐ Missing ☐ Partial ☐ Replacement Ephemeris

Priority* (See below)

Data type(s)* (Select all that apply)

FDS ☐ Carry-out & APID1 ☐ Attitude ☐

Low Rate ☐ 1 ☐ 2 ☐ 3 ☐ 4 ☐ 5 ☐ 6

MODIS ☐ 64

ASTER ☐ EDS ☐ 257/259 ☐ 261/263 ☐ 265/267 ☐ 273/275 ☐ 277/279 ☐ 281/283

289/291 ☐ 293/295 ☐ 297/299 ☐ 305/307 ☐ 309/311 ☐ 313/315

CERES ☐ 131 ☐ 132 ☐ 133 ☐ 167 ☐ 168 ☐ 169

MOPITT ☐ 192 ☐ 193 ☐ 194 ☐ 195 ☐ 196

MISR ☐ 320 ☐ 323 ☐ 325 ☐ 326 ☐ 329 ☐ 330 ☐ 332 ☐ 335 ☐ 337

338 ☐ 340 ☐ 343 ☐ 344 ☐ 347 ☐ 349 ☐ 350 ☐ 353 ☐ 354

356 ☐ 359 ☐ 360 ☐ 365 ☐ 373 ☐ 374 ☐ 378 ☐ 380

Comments

Actionee Org* ☐ EDOS ☐ FOT

* Required field

Data Reorder Priorities

Priority 1: Critical Data needed within 24-48 hours. This includes expedited data, calibration data, and any other data identified by the instrument team as critical path data.

Priority 2: Important Data observation or activity (e.g., targets of opportunity)

Priority 3: All other data needs

[Return to Home Page](#)

Figure 23. Add New Request page for EOS Terra Data Reorder

Some data providers (e.g., Landsat-7) have decided not to support data re-supply; consult appropriate Interface Control Documents to determine suppliers able to re-supply data. The re-ingest may entail ingest procedures that are addressed in a separate lesson on Ingest, or procedures unique to ingest of replacement data (e.g., when replacement data from EDOS are provided on a tape with multiple files, only some of which are needed to replace a lost granule). As with re-processed data, the resultant recovered file will have a new UR.

The following procedure addresses re-ingest of lost data.

Re-Ingest of Lost Data

- 1 Identify the source for each of the lost granules that were ingested.
- 2 If you have not already done so, retrieve the file location metadata for each file (see Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project Procedure 17.6.2.5.1 SDSRV Retrieval of File Location Metadata*).

Note: Except for Step 48, the remaining steps may be performed by the Ingest Technician.

- 3 With reference to the applicable Interface Control Document (ICD) and using the granule metadata retrieved in Step 2, initiate the required data re-supply requests as defined in the ICD for those data suppliers able to re-supply data.
- To re-order Level 0 production data sets (PDSs) from EDOS, the DAACS use the ESDIS-sponsored EOS Data Reorder Web Tool. Steps 4 - 26 address use of the tool to submit a re-order request.
 - EDOS furnishes L0 replacement data to the GSFC Earth Sciences (GES) DAAC on DTF-2 tapes. A tape may contain multiple granules and files, a subset of which are needed to replace the lost granule(s). Steps 27 - 54 address recovery of the lost data.
 - EDOS furnishes L0 replacement data to other DAACs that do not have DTF-2 tape drives. In this case, EDOS transfers the necessary PDS(s) to ECS in the automated Ingest process for polling with delivery record, monitored by the Ingest Technician (see Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project* **Procedure 16.2.5 Monitor/Control Ingest Requests**). The re-order is accomplished as specified in steps 4 - 26; the only other necessary steps in this procedure are 45, 48, and 49.
 - Note: Some data suppliers (e.g., Landsat-7) do not support re-supply of data.

Note: Steps 4 – 8 are to access and launch browser software.

- 4 At the UNIX command shell prompt, type **setenv DISPLAY *clientname*:0.0** and then press the **Return/Enter** key.
- For *clientname*, use either the local terminal/workstation IP address or its machine name.
- 5 Start the log-in to a Netscape host by typing **/tools/bin/ssh *hostname*** (e.g., g0ins02, e0ins02, l0ins02, n0ins02) at the UNIX command shell prompt, and press the **Return/Enter** key.
- If you receive the message, **Host key not found from the list of known hosts. Are you sure you want to continue connecting (yes/no)?** type **yes** (“y” alone does not work).
 - If you have previously set up a secure shell passphrase and executed **sshremote**, a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears; continue with Step 6.
 - If you have not previously set up a secure shell passphrase; go to Step 7.
- 6 If a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears, type your **Passphrase** and then press the **Return/Enter** key. Go to Step 8.
- 7 At the **<user@remotehost>'s password:** prompt, type your **Password** and then press the **Return/Enter** key.
- You are logged in and a UNIX command shell prompt is displayed.

- 8 Type **netscape &** and then press the **Return/Enter** key.
 - The Netscape web browser is displayed.
- 9 Click in the **Netsite:** or **Location:** field.
 - The field is highlighted.
- 10 Type the Universal Resource Locator (URL) for the End-to-End Data Tracking System (<http://edts1.gsfc.nasa.gov:8080/index.html>) and then press the **Return/Enter** key.
 - The **End-to-End Data Tracking System** index page is displayed, offering access to various reports and related links.
- 11 Under **Related Links**, click on the appropriate Data Reorder link.
 - There are links for various satellite platforms (e.g., Terra, Aqua).
 - The appropriate **Data Reorder Request** page for the selected Data Reorder link is displayed.
- 12 On the Data Reorder Request page, click on the **Add New Request** link.
 - The **Add New Request** page is displayed.
- 13 On the **Add New Request** page, click in the **Requestor's Name** field.
 - The cursor is displayed in the field.
- 14 Type the name of the person making the request.
 - The typed entry is displayed in the field.
- 15 Single-click on the pull-down arrow at the end of the **Requestor's Organization** field and then single-click on the name of the requesting organization to select it, or, if the requesting organization is not displayed in the pull-down menu, select **Other**.
 - The choices are **LaRD DAAC**, **LaTIS**, **GSFC DAAC**, **EDC**, **EDOS**, **FOT**, **ESDIS**, **ASTER GDS**, and **Other**.
 - The selected choice is displayed in the field.
- 16 In the **Requested Data (UTC)** block of the page, single-click on the pull-down arrow at the end of the **Year** field and then single-click on the year for the missing data to select it.
 - The selected choice is displayed in the field.
- 17 In the **Requested Data (UTC)** block of the page, single-click on the pull-down arrow at the end of the **DOY** (Day of Year) field and then single-click on the day of the year for the missing data, first scrolling with the scroll bar if necessary to display the desired day.
 - The selected choice is displayed in the field.

- 18** In the **Requested Data (UTC)** block of the page, single-click on the pull-down arrow at the end of the **Start time** field and then single-click on the hour representing the start of a two-hour time window for the missing data.
- The choices are in even two-hour time intervals beginning with **0000** and proceeding to **2200**.
 - The selected choice is displayed in the field.
- 19** In the **Requested Data (UTC)** block of the page, single-click on the pull-down arrow at the end of the **Stop time** field and then single-click on the hour representing the end of a two-hour time window for the missing data.
- The choices are in even two-hour time intervals beginning with **0200** and proceeding to **2400**.
 - The selected choice is displayed in the field.
- 20** In the **Request Reason** block of the page, click on the appropriate radio button to indicate the reason for the request, specifying that the dataset is **Missing** or **Partial**.
- The selected button is filled to indicate its selection.
- 21** For any data priority other than **3**, click on the pull-down arrow at the end of the **Priority** field and then click on the appropriate priority for the request.
- The choices are **1** (critical data needed within 24-48 hours), **2** (important data observation or activity, such as a target of opportunity), and **3** (all other data needs). The default is **3**, and this step may be skipped if that is the priority for the request.
 - The selected choice is displayed in the field.
- 22** Click in the applicable check boxes in the **Data type(s)** block of the page to specify the desired **FDS** or Flight Dynamics System information (Carry-out and APID1 or satellite housekeeping data, Attitude, and Replacement Ephemeris), **Low Rate** information, and instrument (e.g., **MODIS**, **ASTER**, **CERES**, **MOPITT**, and **MISR**) application process identifiers (APIDs).
- The selected check boxes each display a checkmark to indicate selection.
- 23** If it is desirable to enter any comments concerning the request, click in the **Comments** field; otherwise, go to Step 25.
- The cursor is displayed in the field.
- 24** Type any comments to be submitted with the request.
- The typed entry is displayed in the field.
- 25** In the **Actionee Org** block of the page, click on the appropriate radio button to identify the actionee for the request, specifying **EDOS** or **FOT**.
- The selected button is filled to indicate the selection (in this case, **EDOS**).

- 26 Click on the **Submit** button at the bottom of the form.
- Following confirmation, the request submittal is acknowledged with a request ID.
- 27 If the replacement data are on a DTF-2 tape from EDOS, load the tape into a DTF-2 drive, using Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project* **Procedure 16.3.3.4 Load a DTF-2 Drive**.
- The tape is loaded.
- 28 Access a terminal window logged in to the appropriate host (e.g., Distribution Server).
- Examples of Distribution Server host names include **e0dis02**, **g0dis02**, **l0dis02**, and **n0dis02**.
- 29 Type **cd path** and then press the **Return/Enter** key.
- **path** represents the directory path to the location where the data from the EDOS archive tape should be copied.
 - Using an empty directory would help identify the data from the tape.
- 30 Type **tar xvf device** and then press the **Return/Enter** key.
- **device** is the DTF-2 drive device name (e.g., **/dev/rmt/2n**) as it is known to the shell.
 - For example:
tar xvf /dev/rmt/2n
 - As files are read from the tape the file names, file sizes (in bytes), and number of blocks are listed on the screen.
 - For example:
x DZ9ZA49.MDR, 17393 bytes, 34 tape blocks
- 31 Type **pg PPMUDR_name** and then press the **Return/Enter** key.
- **PPMUDR_name** represents the file name of the PDS Physical Media Unit Delivery Record (PPMUDR).
 - The PPMUDR file name has a **.MDR** extension.
 - The PPMUDR is the first item on the EDOS archive tape.
 - For example:
pg DZ9ZA49.MDR
 - Although this procedure has been written for the **pg** command, any UNIX editor or visualizing command (e.g., **vi**, **view**, **more**) can be used to review the log file.

32 Observe the contents of the PPMUDR to identify the PDS(s) to be archived.

- Packet date/time ranges in the PPMUDR can be used to determine which PDS(s) is (are) to be archived.
 - In the PPMUDR the PDSs on the tape are listed in file groups, which represent data sets [i.e., science data file(s) and corresponding metadata file].
 - Each file group (data set) includes the date/time range of the data specified as FIRST_PACKET_TIME and LAST_PACKET_TIME.
- For example (extract from a PPMUDR):

OBJECT = FILE_GROUP

DATA_SET_ID = P0420064AAAAAAAAAAAAAAAAA03101231459600

DATA_TYPE = MOD000

FIRST_PACKET_TIME = 2003-04-10T00:00:00.000000Z

LAST_PACKET_TIME = 2003-04-10T01:59:59.999999Z

PACKET_COUNT = NOT USED

OCTET_COUNT = NOT USED

TEST_FLAG = F

APID_COUNT = 1

OBJECT = APID_SPEC

APID_IN_PDS = 64

END_OBJECT = APID_SPEC

FILE_COUNT = 2

OBJECT = FILE_SPEC

DIRECTORY_ID = NOT USED

FILE_ID = P0420064AAAAAAAAAAAAAAAAA03101231459600.PDS

FILE_TYPE = METADATA

FILE_SIZE = 384

END_OBJECT = FILE_SPEC

OBJECT = FILE_SPEC

DIRECTORY_ID = NOT USED

FILE_ID = P0420064AAAAAAAAAAAAAAAAA03101231459601.PDS

FILE_TYPE = DATA

FILE_SIZE = 108000

END_OBJECT = FILE_SPEC

END_OBJECT = FILE_GROUP

- In the preceding example one data set is defined (as a “FILE_GROUP”).
 - The data type for the set is MOD000.
 - The data were collected on April 10, 2003 between midnight GMT (00:00:00.000000Z) (FIRST_PACKET_TIME) and just before 2:00 A.M. GMT (01:59:59.999999Z) (LAST_PACKET_TIME).
 - There are two files in the data set (FILE_COUNT = 2).
 - One file (P0420064AAAAAAAAAAAAAAAA03101231459600.PDS) is a metadata file (in EDOS terminology, a “construction record”).
 - The other file (P0420064AAAAAAAAAAAAAAAA03101231459601.PDS) is a data file.
 - Based on information embedded in the file names, the data set was created on April 11, 2003 at 11:14:59 P.M. (as described under the next bullet).
- The EDOS archive tape may contain both nominal and reprocessed PDSs but creation times in file names differentiate between the versions.
 - Ingest the latest (most recent) version if there is more than one version.
 - PDS file names consist of 40 bytes (characters) and Bytes 23 through 33 specify the creation time of the file.
 - For example, **03101231459** is the creation time in the following file name:
P0420064AAAAAAAAAAAAAAAA03101231459601.PDS
 - **03** indicates the year (2003).
 - **101** specifies the Julian day (April 11, the 101st day of the year).
 - **231459** is the time of file creation (11:14:59 P.M.).
- It is the Archive Manager’s responsibility to resolve any questions concerning which PDSs should be archived (see step 2).

33 Type **cp filename1 filename2 [... filenameN] path** and then press the **Return/Enter** key.

- **filename1 filename2 [... filenameN]** represent the file names of the PDS files to be ingested.

- Copy both the data and metadata files (as identified in the PPMUDR) for each data set.
- ***path*** is the directory path to the Ingest polledEDOS directory; i.e., the directory in which the ECS software for EDOS ingest routinely looks for EDOS delivery records and data.
 - The EDOS polling directory is specified as a parameter in the Registry database or in the configuration file for EDOS polling (e.g., EcInPolling.EDOS.CFG).
- For example:


```
cp P0420064AAAAAAAAAAAAAAAA03101231459600.PDS
P0420064AAAAAAAAAAAAAAAA03101231459601.PDS
/usr/ecs/OPS/CUSTOM/icl/x0icg01/data/pollEDOS
```
- **NOTE:** If a DAAC-unique script is available for creating delivery records and signal files and placing the files in the polling directory, use the script and skip Steps 34 through 44 (go to Step 45 after running the script). Otherwise, manually generate delivery records and signal files as described in Steps 34 through 44.

34 Type **cd *path*** and then press the **Return/Enter** key.

- ***path*** is the directory path to the Ingest polledEDOS directory.
- For example:


```
cd /usr/ecs/OPS/CUSTOM/icl/x0icg01/data/pollEDOS
```

NOTE: Steps 35 through 39 describe how to use an old delivery record (PDR) as a template for generating a new PDR.

35 Type **cp *old_PDR_filename new_PDR_filename*** and then press the **Return/Enter** key.

- ***old_PDR_filename*** represents the file name of an old PDR that is being used as a template for creating a PDR for PDS files to be ingested.
- ***new_PDR_filename*** represents the file name of the new PDR that is being created for PDS files to be ingested.
 - Use the EDOS file-naming convention for PDRs (refer to the EDOS ICD, 423-ICD-EDOS/EGS):
 - PDR file names consist of 38 bytes (characters).
 - Byte 1 identifies the file as either a PDS Delivery Record (“X”) or EDS Delivery Record (“Y”).
 - Bytes 2 through 8 identify the spacecraft ID (SCID) (three bytes) and first Applications Process Identifier (APID) (four bytes) in the data set (right-justified and, if necessary, zero-filled on the left).

- Bytes 9 through 15 identify the SCID and second APID in the data set (right-justified and, if necessary, zero-filled on the left), if applicable. If no second APID is present in the data set, this item has a value of “AAAAAAA”.
- Bytes 16 through 22 identify the SCID and third APID in the data set (right-justified and, if necessary, zero-filled on the left), if applicable. If no second APID is present in the data set, this item has a value of “AAAAAAA”.
- Bytes 23 through 33 identify the GMT/ZULU time when the data set was created.
- Byte 34 is a numeric identification in the range of “0” to “9” to aid in distinguishing the order of data set creation during the day and to provide uniqueness to the file name.
- Bytes 35 through 38 are the file name extension (i.e., “.PDR” or “.EDR”)
- For example:

X0420064AAAAAAAAAAAAAAAA031012314596.PDR

X identifies the file as a PDS Delivery Record.

0420064 identifies the SCID (**042** = Terra) and first APID (**0064** = MOD000 data type) in the data set.

AAAAAAAA indicates that there is no second APID in the data set.

AAAAAAAA indicates that there is no third APID in the data set.

03101231459 is the GMT/ZULU time when the data set was created [**03** indicates the year (2003); **101** specifies the Julian day (April 11, the 101st day of the year); **231459** is the time of data set creation (11:14:59 P.M.)].

6 is a numeric identifier (sixth data set of the day).

.PDR is the file-name extension for a PDS Delivery Record.

36 Type **vi new_PDR_filename** and then press the **Return/Enter** key.

- The PDR template file is opened (displayed by the vi text editor).
- Although this procedure has been written for the **vi** editor, any UNIX editor can be used to create the PDR.

37 Using vi editor commands modify the PDR file to specify ingest of one of the data sets to be ingested.

- Create a separate PDR for each data set [science data file(s) and corresponding metadata file – refer to the PPUDR “file group” example in Step 32].

- The following vi editor commands are useful:

- **h** (move cursor left).
- **j** (move cursor down).
- **k** (move cursor up).
- **l** (move cursor right).
- **a** (append text).
- **i** (insert text).
- **r** (replace single character).
- **x** (delete a character).
- **dw** (delete a word).
- **dd** (delete a line).
- **ndd** (delete *n* lines).
- **u** (undo previous change).
- **Esc** (switch to command mode).

38 Press the **Esc** key.

39 Type **ZZ**.

- New PDR file is saved.
- UNIX prompt is displayed.

40 Type **vi XFR_filename** and then press the **Return/Enter** key.

- A new file with the specified **XFR_filename** is opened.
 - Use the EDOS file-naming convention for signal files (refer to the EDOS ICD, 423-ICD-EDOS/EGS):
 - Signal file name is the corresponding PDR file name plus the signal file name extension (i.e., “.XFR”).
 - For example:

X0420064AAAAAAAAAAAAAAAA031012314596.PDR.XFR

- The signal file indicates that the relevant data files and PDR have been put in the polling directory and are ready to be ingested.
- Although this procedure has been written for the **vi** editor, any UNIX editor can be used to create the signal file.

41 Using vi editor commands create a file that contains the name of the relevant PDR.

- A signal file contains the name of the relevant PDR only.
- For example:

X0420064AAAAAAAAAAAAAAAA031012314596.PDR.

42 Press the **Esc** key.

43 Type **ZZ**.

- New signal file is saved.
- UNIX prompt is displayed.
- At the next polling occasion, the EDOS polling client should detect the signal file and initiate ingest of the data specified in the corresponding PDR.

44 Repeat Steps 34 through 43 as required to create delivery records and signal files for all remaining data sets (from the EDOS archive tape) to be ingested.

45 To monitor Ingest request processing (polling with delivery record), perform the procedure for **Monitoring/Controlling Ingest Requests** (see Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project Procedure 16.2.5 Monitor/Control Ingest Requests*).

46 Remove the EDOS-provided tape from the DTF-2 drive, using Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project Procedure 16.3.3.5 Unload a DTF-2 Drive*.

47 Verify that the data have been inserted into the archive as described in the procedure for **Verifying the Archiving of Ingested Data** (see Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project Procedure 16.2.10 Verify the Archiving of Ingested Data*).

48 When insertion into the archive has been verified, the Archive Manager specifies "set delete" for the replaced data/metadata by using procedures for granule deletion to mark the data/metadata for deletion from the archive (see Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project Section 17.4.2 Selecting Granules for Deletion and Procedure 17.4.2.3 Selection Using a Separate Input File*).

49 When insertion into the archive has been verified, ensure that the EDOS archive tape is returned to the EDOS Level 0 Processing Facility (LZPF).

NOTE: Clean up (as described in Steps 50 through 54) the directory into which data were originally copied from the EDOS archive tape. If preferred, skip Steps 50 through 54 and use the script described in the procedure **Clean the Polling Directories** (see Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project Procedure 16.2.11 Clean the Polling Directories*).

- 50 Type **cd *path*** and then press the **Return/Enter** key.
- *path* represents the directory path to the location where the data from the EDOS archive tape were first copied.
- 51 Type **ls** and then press the **Return/Enter** key.
- A listing of the files in the current directory is displayed.
- 52 Type **rm *filename1 filename2 [... filenameN]*** and then press the **Return/Enter** key.
- *filename1 filename2 [... filenameN]* represent the names of the files to be removed from the directory.
 - A wildcard may be used if some of the files have common characteristics.
 - For example:
rm *.PDS
 - A prompt is displayed requesting whether or not a particular file should be removed.
 - For example:
rm: remove DZ9ZA49.MDR (yes/no)?
- 53 Type **y** and then press the **Return/Enter** key.
- The specified file is deleted and (if applicable) a prompt is displayed requesting whether or not another particular file should be removed.
- 54 Repeat Step 53 as necessary.

Restoring the AMASS Database

The AMASS database is restored manually by the System Administrator or the Archive Manager using the AMASS command **amassrestore**. This command restores the last full backup, the last partial backup, and all journal transactions that have occurred since the last backup. It creates a sub-directory under filesysdb called **journal**. All restored files are copied to the **journal** directory. The following restore procedure uses a backup volume or tape device.

Restore AMASS Database

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01) as **amass** or **root**.

Caution

Do not use the **amassrestore** command when AMASS is running. To shutdown AMASS, refer to the Special Shutdown Procedures in the AMASS technical documentation *Installing AMASS*.

- 2 To inactivate the AMASS file system, type **/usr/amass/bin/amassstat -i**.
 - The AMASS file system is inactivated.
- 3 Make sure the backup drive is available.
 - If there is another volume in the drive, return it to its home slot by entering **/usr/amass/daemons/amassrecovery -s** (the option **-s** prevents system startup and performs file recovery).
- 4 Type **/usr/amass/bin/amassrestore -v -L <barcodelabel>** and then press the **Return/Enter** key.
 - If you do not know the barcode label number for the backup volume, it can be obtained by entering **/usr/amass/bin/vollist 1**.
 - The AMASS database is restored from the backup volume.

Restoring the ACSLS Database

ACSLs provides the **rdb.acsss** utility to restore the database in case of severe disk or data problems. If you have made regular backups, it should be possible to restore the database with little or no loss of data. Restoring the database is likely to be necessary if there has been a system crash, or if the database can not be started or has a physical or logical error. The following procedure is applicable.

Restore the ACSLS Database

- 1 At the host for ACSLS (e.g., **e0drs03**, **g0drs03**, **l0drs02**, **n0drs03**), log in using the **acsss** user ID and password.
 - The **acsss** command-process window is displayed with the **ACSSS>** prompt.
- 2 Load the restore tape into the backup drive.

3 Type **rdb.acsss**, and then press the **Return/Enter** key.

- If you enter **bdb.acsss** with no options, the backup utility defaults to the default tape device attached and configured to the ACSLS server.
- The system displays the following message.

```
Check tape device (/dev/rmt/0) to make sure you have a tape in the
tape drive.
```

```
[ Hit RETURN to continue or Ctrl-C to exit ]
```

4 Press the **Return/Enter** key.

- The rdb.acsss utility restores the ACSLS database and miscellaneous library resource files.
-

This page intentionally left blank.

AMASS Graphical User Interface

AMASS offers a Graphical User Interface (GUI) called the AMASS Administration Window (AAWin) through which operators can administer volumes and volume groups that are managed by AMASS. AAWin provides a point-and-click interface for identifying volumes their groups, and their configurable parameters. Figure 24 shows the AAWin main window, which is composed of a menu bar, a large middle section called the *workroom*, a utility bar at the right with icons for a trash can, a volume group, and a volume, and a status bar at the bottom with indicator “lights” that represent the current status of AMASS. The figure shows how the window looks when the volume icon on the utility bar has been selected to populate the workroom with icons for volumes, and illustrates the type of volume-related information that appears in a pop-up display as the cursor is moved over one of the icons.

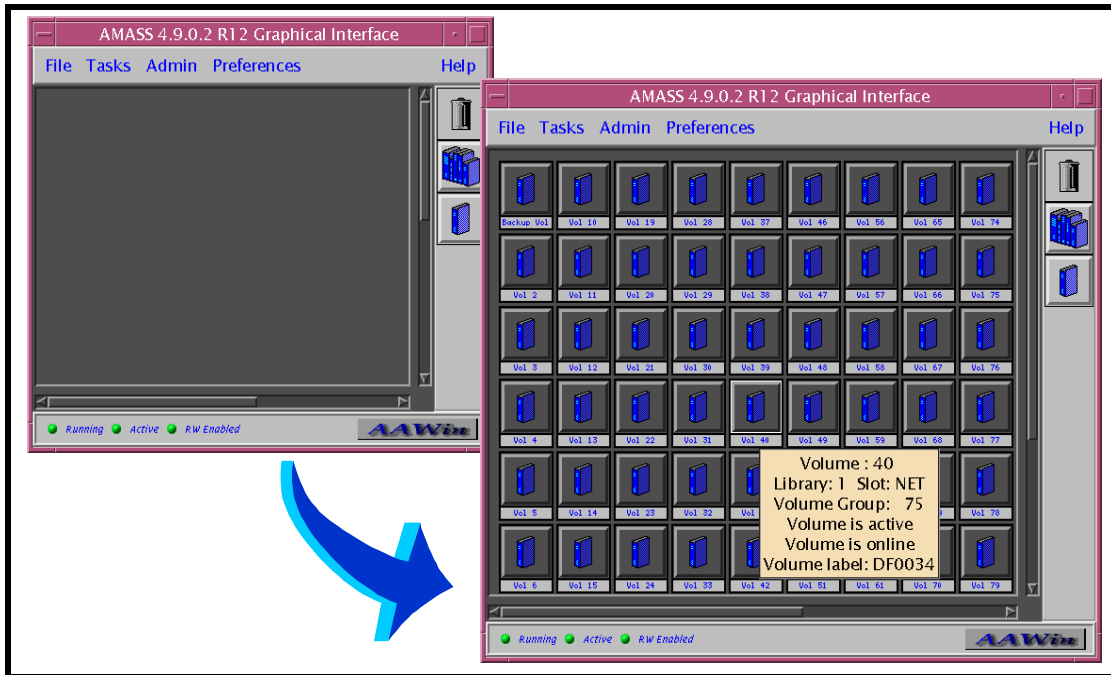


Figure 24. AMASS AAWin Main Window

Use the following procedure to launch the AMASS GUI and view information about volume groups and volumes in the archive.

Using the AMASS GUI to View Volume Group and Volume Information

- 1 Log in as **amass** at the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01).
 - 2 Type **/usr/amass/bin/aawin** and then press the **Return/Enter** key.
 - The AMASS GUI main window is displayed.
 - 3 Click on the **View by Volume Groups** button (middle button at the right of the *workroom*).
 - The *workroom* is populated by icons for volume groups.
 - The **Block List** window is displayed; it is a vertically scrolled list of blocks of items (in this case, volume groups). The *workroom* can display up to 256 icons; the **Block List** window provides access to additional items in blocks of 256.
 - 4 Move the cursor over one of the icons for a volume group.
 - A pop-up display shows data for the volume group (**Volume Group, Volumes in Group, Free Space, Dead Space, Error Count**).
 - 5 Click on the **View by Volumes** button (at the bottom right side of the *workroom*).
 - The *workroom* is populated by icons for volumes.
 - The **Block List** window is also displayed; it is a vertically scrolled list of blocks of items (in this case, volumes).
 - 6 Move the cursor over one of the icons for a volume.
 - A pop-up display shows data for the volume group (**Volume, Library, Slot, Volume Group, Volume Status, Volume Label**).
-

Modify a Volume Group

Figure 25 shows the *Modify a VG* window. This window is opened by selecting *Modify a Volume Group* from the *Tasks* menu. The window is used to modify the characteristics of a volume group. The top portion of the window (not modifiable) lists root directories already configured for a volume group. The middle portion of the window permits adding directories to the list of root directories for the specified volume group. The third major portion of the window, near the bottom, contains indicators of the status of the volume group and buttons for selecting a volume group, as well as buttons across the very bottom of the window for accepting or canceling the modifications. (**Note:** The *Modify a VG* window also is opened if you have the *workroom* populated with volume group icons and you click on one of them. However, in this case you may only modify the volume group on which you clicked; the bottom of the window will not display buttons for selecting a volume group.)

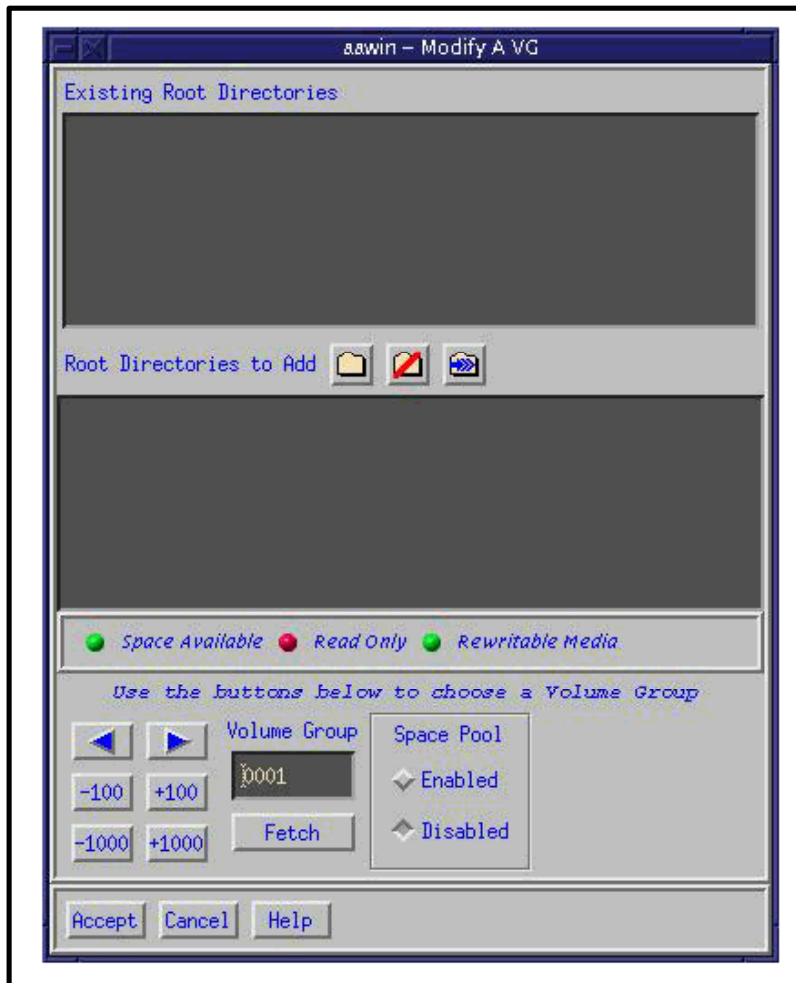


Figure 25. Modify a VG Screen of the AMASS GUI

To see how a volume group can be modified, let's examine how you might assign a new root directory in the AMASS file system to a volume group. Use the following procedure to create the directory and then modify a volume group.

Modify a Volume Group

- 1 Open a second terminal window (other than the one used to launch the AMASS GUI).
 - **NOTE:** This procedure assumes that the AMASS GUI is open after previously being launched from a terminal window.
- 2 In the second terminal window, log in as **amass** at the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01).
- 3 To change to the **dss_amass** directory, type **cd /dss_amass**, and then press the **Return/Enter** key.
 - The working directory is changed to **/dss_amass**.

- 4 To create an empty directory with path **/dss_amass/training/** to assign to the volume group, type **mkdir training**, and then press the **Return/Enter** key.
- 5 On the AMASS GUI main window, click on the **View by Volume Groups** button (middle button at the right of the *workroom*).
 - The *workroom* is populated by icons for volume groups.
 - The **Block List** window is also displayed; it is a vertically scrolled list of blocks of items (in this case, volume groups).
- 6 Follow menu path **Tasks→Modify a Volume Group**.
 - The **Modify a VG** window is displayed, showing data for Volume Group 0001.
- 7 In the area for choosing a volume group, near the bottom of the window, use the buttons to set the number displayed in the **Volume Group** field to the desired volume group.
 - A click on the right-pointing arrow button or the left-pointing arrow button respectively increases or decreases the number by one. Buttons below the arrow buttons may be used to increase or decrease the number in multiples of 100 or 1000, as indicated on the buttons.
- 8 When the **Volume Group** field displays the number of the desired volume group, click on the **Fetch** button.
 - The list of root directories already configured for the selected volume group is displayed in the **Existing Root Directories** field.
 - The status indicators show the status of the selected volume group.
- 9 Click on the **File/Directory Selection** button (leftmost button after the label **Root Directories to Add**, with folder icon).
 - A **File Selection** filter window is displayed.
- 10 In the **File Selection** filter window, click on the **Filter** button.
 - The **Filter** field displays **/usr/amass/***, and directories and files are displayed in the **Directories** and **Files** windows, respectively.
- 11 Use the **Filter** button and selection of directories in the **Directories** window to display **/dss_amass/training/** in the **Selection** field.
 - The **Selection** field displays **/dss_amass/training/**.
- 12 In the **File Selection** filter window, click the **OK** button.
 - The **Root Directories to Add** field of the **Modify a VG** window displays **/dss_amass/training/**.

- 13 To examine the capability to edit the list of Directories to Add, click on the entry **/dss_amass/training/** to highlight it in the **Root Directories to Add**, then click on the **Remove a File/Directory from List** button (middle button after the label **Root Directories to Add**, with folder icon crossed out with a red line).
 - The entry **/dss_amass/training/** is removed from the **Root Directories to Add** field.
 - 14 Repeat steps 10 - 12 to restore the entry **/dss_amass/training/** to the **Root Directories to Add** field.
 - The **Root Directories to Add** field of the **Modify a VG** window displays **/dss_amass/training/**.
 - 15 In the **Modify a VG** window, click on the **Accept** button at the bottom of the window.
 - The entry **/dss_amass/training/** is removed from the **Root Directories to Add** field and appears in the **Existing Root Directories** field.
 - The **Modify a VG** window is closed.
-

Modify a Volume

Figure 26 shows the *Modify a Volume* window. This window is opened by selecting *Modify a Volume* from the *Tasks* menu. The window is used to modify the characteristics of a volume. The right side of the window shows the current set of statistics and configuration information (not modifiable) for the volume listed in the **Volume** field on the left side of the window (the **Volume** field looks like a button, but if you click on it, a “spinbox” is displayed, with arrow buttons permitting increases or decreases to the volume number, and buttons at the bottom to **Accept** or **Cancel** the change; accepting the change closes the spinbox, displays the new number in the **Volume** field, and displays data for that volume). The left side of the *Modify a Volume* window provides access to modifiable characteristics of the volume. Changes made to the buttons and fields in the window do not take effect until the **Accept** button at the bottom of the window is clicked. (**Note:** The *Modify a Volume* window also is opened if you have the *workroom* populated with volume icons and you click on one of them. However, in this case you may only modify the volume on which you clicked; the **Volume** field does not look like a button and may not be changed.)

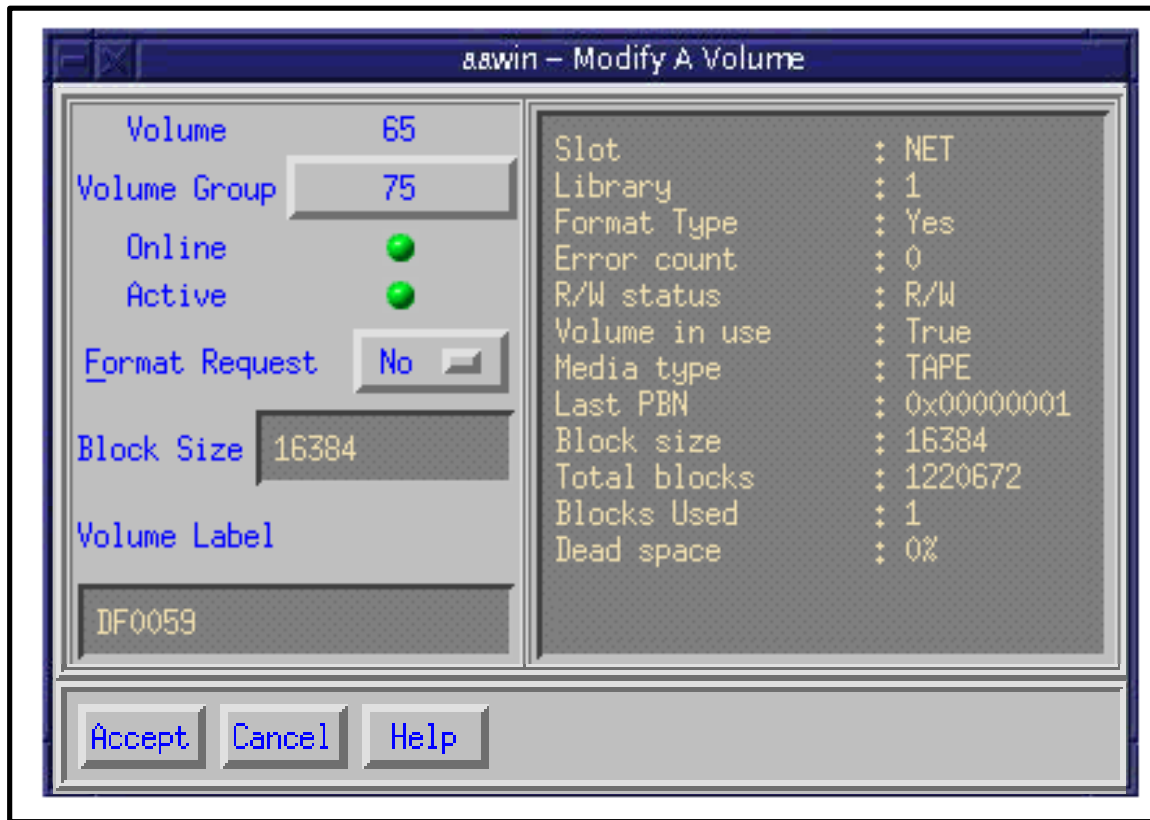


Figure 26. Modify a Volume Window of the AMASS GUI

There are six fields that can be edited for a given volume:

1. The first is a button for setting the **Volume Group**. Clicking the button opens a spinbox for selecting the volume group to which the volume is to be assigned.
2. Below the Volume Group button is an **Online/Offline** indicator light with label. Clicking on the indicator toggles its state and updates the text field (label) next to it.
3. Below the Online/Offline indicator is an **Active/Inactive** indicator light with label. Clicking on the indicator toggles its state and updates the text field (label) next to it.
4. Next is a **Format Request** option button permitting selection of a formatting option for the volume.
5. Next is the **Block Size** field, applicable only to tape libraries when a format is requested to be done on the volume. This field requires a numeric value, which should be a multiple of 16384.
6. The last modifiable field is a text field for specifying the volume label.

To examine the functioning of the *Modify a Volume* window, use the following procedure.

Modify a Volume

- 1 On the AMASS GUI main window, click on the **View by Volume Group** button (middle button at the right of the *workroom*).
 - The *workroom* is populated by icons for volume groups.
 - The **Block List** window is also displayed; it is a vertically scrolled list of blocks of items (in this case, volume groups).
- 2 Click on the icon for a high-numbered volume group with several volumes in it.
 - The *workroom* is populated with icons for the volumes in the selected volume group, and the **Modify a VG** window is displayed, showing data for the selected volume group.
- 3 Click on the icon for one of the volumes.
 - The **Modify a VG** window is closed and the **Modify a Volume** window is displayed, showing data for the selected volume.
- 4 To change the volume group to which the volume is to be assigned, note the **Volume Group** number indicated on the **Volume Group** button, and then click on the button..
 - A spinbox is displayed showing the **Volume Group** number, with right-pointing and left-pointing arrow buttons respectively to increase or decrease the number.
- 5 Use the arrow buttons to change the **Volume Group** number, and then click on the **Accept** button in the spinbox.
 - The spinbox is closed and the new number appears in the **Modify a Volume** window as the **Volume Group** number.
- 6 Return the **Volume Group** number to its original value by repeating steps 4 and 5, using the spinbox to set the number to that which you noted originally.
 - The spinbox is closed and the original number appears in the **Modify a Volume** window as the **Volume Group** number.
- 7 Experiment with the **Online/Offline** and **Active/Inactive** indicators.
 - When the volume is indicated to be **Online**, clicks on the **Active/Inactive** indicator toggle the color and label for the indicator.
 - When the volume is indicated to be **Inactive**, clicks on the **Online/Offline** indicator toggle the color and label for the indicator.
- 8 Click on the **Format Request** option button.
 - A pop-up option menu is displayed for selection of **Yes** or **No**, and when one of those options is clicked, the indicated choice is displayed on the option button.

- 9 Use the mouse to move the cursor to the **Block Size** field.
 - A blinking cursor appears in the **Block Size** field.
 - 10 Use the keyboard to enter or change the value in the **Block Size** field.
 - The entered data appear in the **Block Size** field.
 - 11 Use the mouse to move the cursor to the **Volume Label** field.
 - A blinking cursor appears in the **Volume Label** field.
 - 12 Use the keyboard to enter or change the value in the **Volume Label** field.
 - The entered data appear in the **Volume Label** field.
 - 13 If you wish to cancel any request for changes to the volume, click on the **Cancel** button at the bottom of the window. If you wish to accept the changes, click on the **Accept** button at the bottom of the window.
 - When you click the **Accept** button, *AAWin* attempts to make the requested changes. For most changes, specifically changes to **Online/Offline** and **Active/Inactive** status, the requested **Volume Group** for the volume, and the **Volume Label**, the changes can be made immediately. But if a format has been requested, then the **Online/Offline** and **Active/Inactive** status changes are not applied immediately. Instead, the requests for these status changes and the format changes are passed to the *AAWin Scheduler* daemon for processing. Changes made by the **Scheduler** occur when the job is processed, which depends on how many other jobs are currently scheduled.
-

Archive Monitoring and Troubleshooting

Previous sections of this lesson have addressed the use of tools that can assist you in monitoring the Archive. Specifically, the System Requests tab of the Science Data Server Operator GUI provides a view into archive request processes, the Data Distribution GUI provides information on distribution requests, and the AMASS GUI offers ready access to information about the status of archive volumes and volume groups. If archive problems arise, there are additional resources that can provide more detailed monitoring and assistance in troubleshooting. Troubleshooting is a process of identifying the source of problems on the basis of observed symptoms. Because the Archive is at the heart of ECS and its Data Server interacts with so many subsystems, problems with Archive functions may be traced to the Data Server subsystem (DSS) or one of many other ECS subsystems, including (but not limited to) the following:

- Ingest Subsystem (INS).
- Planning Subsystem (PLS).
- Data Processing Subsystem (DPS).
- Interoperability Subsystem (IOS).
- Communications Subsystem (CSS).
- Data Management Subsystem (DMS).

Table 3 summarizes actions to be taken in response to some common Archive problems. If the problem cannot be identified and fixed without help within a reasonable period of time, the appropriate response is to call the Help Desk and submit a trouble ticket in accordance with site Problem Management policy.

Table 3. Troubleshooting Archive Problems

Symptom/Problem	Response
Unable to log in to the FSMS host (e.g., x0drg01)	Check with the Operations Controller/System Administrator to ensure that the host is "up."
AMASS is not running	Have the System Administrator restart AMASS.
A volume is inactivated by AMASS	Check for AMASS errors and, unless there are many errors, use the command <code>/usr/amass/bin/volstat -a <vol_number></code> to re-activate the volume. [For detailed instructions, refer to the procedure Use the amass_log script to Display AMASS Errors (subsequent section of this lesson).]
A storage system robot gets out of synchrony with AMASS concerning the location of media.	Re-establish synchrony. [For detailed instructions, follow the procedure Use mediamove to Establish Synchrony Between quedisplay and medialist (subsequent section of this lesson).]
An Ingest or Data Processing action cannot complete because of failure to store data (reflected as failure on Ingest or Processing GUIs).	Check to ensure AMASS is on line; check for file copy errors, network problems, mount point problems. [For detailed instructions, refer to the procedure for Recovery from Failure to Store or Retrieve Data (subsequent section of this lesson).]

Although there are no custom ECS reports of archive events, the Storage Management GUIs provide a **Storage Events** tab, as shown in Figure 27. This tab allows an operator to review events in the storage management Event Log. The Event Log Search Parameters box permits specification of constraints to limit the event log search that is executed when the **Search** button is clicked. This box enables an operator to select a date interval for the search, and it provides option buttons to specify the type and level of events to be displayed. Options on the **Event Type** option button are: **Any**, **Device**, **Cache**, **Software**, **COTS**, **Sybase**, **Pulldisk**, and **Unknown**. Options on the **Event Level** option button are: **Any**, **Information**, **Warning**, **Error**, **Severe**, **Fatal**, and **Unknown**. The resulting events are then displayed in the **Event Log** field. A **Purge Selected** button at the bottom of the window permits deletion of entries in the Event Log that have been selected, or highlighted, by clicking on them in the **Event Log** field.

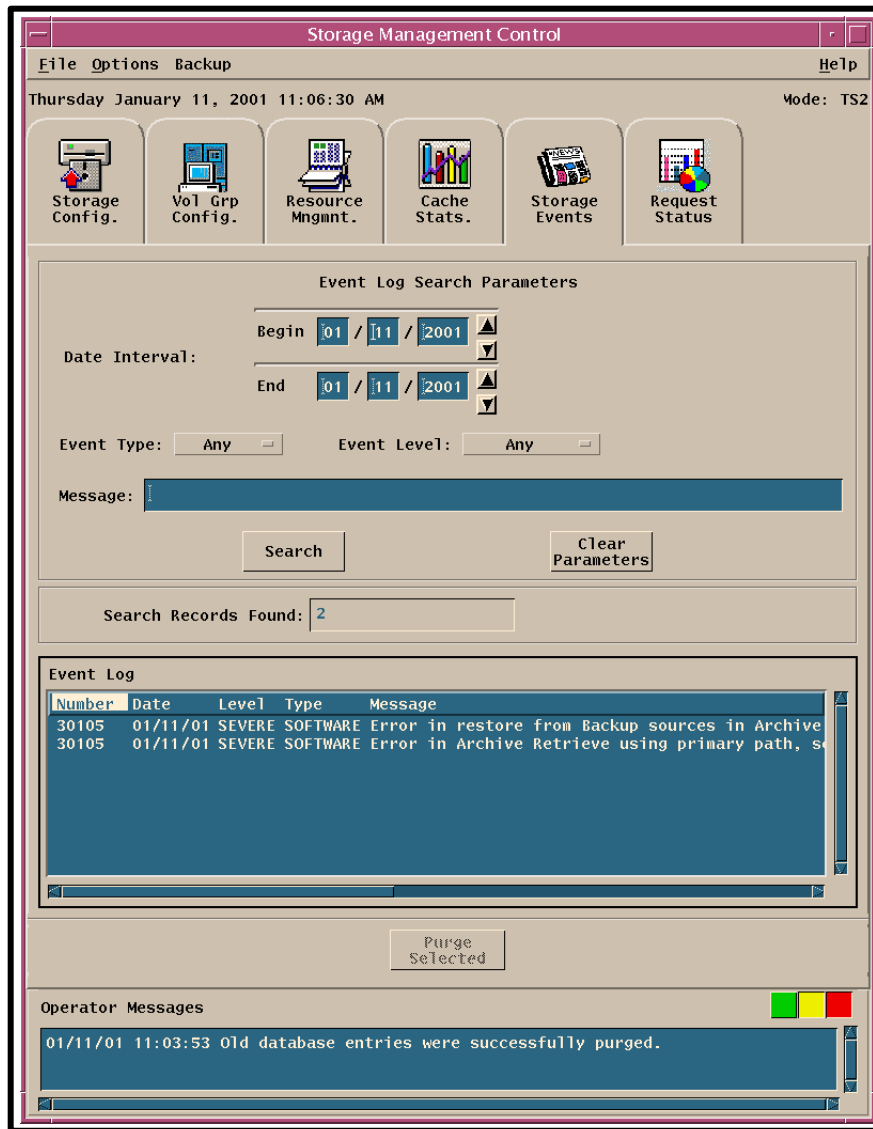


Figure 27. Storage Management, Storage Events Tab

The Storage Management GUI also provides a powerful capability for tracking the status of requests associated with archive activities (see Figure 16 and associated procedure, page 42). This is the **Request Status** tab, which permits filtering requests to specific Storage Management request types and quick determination of the status of an individual request. Identification of the status of an archive request using this tool can be a helpful first step in troubleshooting apparent archive problems.

There are also several troubleshooting tools provided with AMASS that can assist you in monitoring archive activity and in responding to fault notifications. The *AMASS System Administrator's Guide* includes instructions on using these tools. Some of the most useful ones are addressed here.

AMASS Commands, Utilities, and Scripts for Monitoring and Fault Response

The AMASS file system needs to have the following daemons running at all times:

- **amassmain.**
- **daemons/lm_ip -a fslock.**
- **qset.**
- **klogd.**
- **amass_iocomp.**
- **libsched** (one instance for each virtual library).
- **libio_tape** (at least one instance for each drive in each jukebox).

The UNIX process search provides an easy check for these daemons. If they are up, the AMASS **healthcheck** command provides a useful check on the health of AMASS while it is running. Checking for these daemons and running **healthcheck** can be an appropriate initial response to an indication of a potential problem with AMASS. Use the following procedure.

Checking Daemons and Using *healthcheck*

- 1 Log in as **amass** at the FSMS host.
- 2 Type **ps -ef | grep amass** and press the **Return/Enter** key.

- UNIX returns running AMASS processes in a format similar to the following:

```
amass      7214464      7208385    0   Sep 19 ?          0:00 libio_tape 2 1
amass      7208385
/usr/amass/daemons/amassmain 0
amass      7214747      7208385    0   Sep 19 ?          0:10 amass_iocomp
amass      7282853      7215637    0   Sep 20 ?          1:47 libio_tape 1 1
amass      7282868      7215637    0   Sep 20 ?          0:00 libio_tape 1 1
amass      6949087      7215637    0   Sep 20 ?          1:47 libio_tape 1 1
amass      7214915      7208385    0   Sep 19 ?          0:00 klogd
amass      7214972      7208385    0   Sep 19 ?          50:54 libio_tape 1 2
amass      5539722      7217884    0   Sep 20 ?          0:23 libio_tape 1 3
amass      7301726      7215964    0   Sep 20 ?          1:05 libio_tape 3 1
amass      7215313
1      0      Sep 19 ?          9:34
/usr/amass/daemons/lm_ip -a fslock1 -u 128 -f 256 -q 128
amass      7357656      7216363    0   Sep 20 ?          0:00 libio_tape 3 3
amass      7215637      7208385    0   Sep 19 ?          84:10 libio_tape 1 1
amass      7215638      7208385    0   Sep 19 ?          2:43 libsched 3
amass      7277545      7214972    0   Sep 20 ?          0:41 libio_tape 1 2
amass      7215870      7208385    0   Sep 19 ?          2:52 libsched 1
amass      7215964      7208385    0   Sep 19 ?          109:25 libio_tape 3 1
amass      7216363      7208385    0   Sep 19 ?          84:16 libio_tape 3 3
amass      6950984      7217884    0   Sep 20 ?          0:23 libio_tape 1 3
amass      8175053      7212410    0   Sep 26 ?          0:00 libio_tape 1 4
amass      7340525      7217134    0   Sep 20 ?          1:19 libio_tape 3 2
amass      7278745      7217884    0   Sep 20 ?          0:23 libio_tape 1 3
amass      7216941      7208385    0   Sep 19 ?          0:32 qset
```

```

amass      7340710      7217134  0   Sep 20 ?      1:19 libio_tape 3 2
amass      7217134      7208385  0   Sep 19 ?      138:26 libio_tape 3 2
amass      7359550      7216363  0   Sep 20 ?      0:52 libio_tape 3 3
amass      7217388      7208385  0   Sep 19 ?      0:00 libio_tape 2 2
amass      7285477      7215637  0   Sep 20 ?      1:47 libio_tape 1 1
amass      7285537      7215637  0   Sep 20 ?      1:47 libio_tape 1 1
amass      7217884      7208385  0   Sep 19 ?      17:37 libio_tape 1 3
amass      7279821      7214972  0   Sep 20 ?      0:41 libio_tape 1 2
amass      6878049      7208385  0   Sep 19 ?      2:36 libsched 2
amass      7279907      7214972  0   Sep 20 ?      0:41 libio_tape 1 2
amass      7335573      7217134  0   Sep 20 ?      1:19 libio_tape 3 2
.
.
.

```

- If the running processes do not include **amassmain**, **daemons/lm_ip -a fslock**, **qset**, **klogd**, **amass_iocomp**, **libsched**, and **libio_tape**, it may be necessary to restart AMASS (refer to procedure **Rebooting AMASS** [previous section of this lesson]).

- 3 To check the AMASS database integrity, check the availability of write resources FNODEs and cache blocks, and to verify cache partitions, type **/usr/amass/bin/healthcheck -viwc** and press the **Return/Enter** key.

- AMASS returns information on its health in format similar to the following:

```

--- STARTING DATABASE INTEGRITY CHECK ---

    -api has been opened properly and AMASS is running.
    -verifying pathnames.
    -got locks on database
    -unlocking database tables and exiting

--- CHECK COMPLETED!! ---

--- CHECKING AVAILABILITY OF WRITE RESOURCES FNODEs AND CACHE
BLOCKS ---

    -api has been opened properly and AMASS is running.
    -Initializing the passed arguments.
    -Returning the passed arguments.
    -Restoring signals.
    -exiting.

--- CHECK COMPLETED!! ---

--- RUNNING CACHE TEST ---

    -api has been opened properly and AMASS is running.
    -Validating the raw cache.
    -Restoring signals.
    -exiting.

--- TEST COMPLETED!! ---

```

- If an error message is returned, it may be necessary to restart AMASS (refer to procedure **Rebooting AMASS** [previous section of this lesson]).

4 To check library components, type `/usr/amass/bin/healthcheck -vl 1 0 volumenumber` and press the **Return/Enter** key.

- The argument **-l** (lower-case l) specifies the library components check, and requires specification of a jukebox (**1** in this case), a drive number (entering **0** as in this case checks all active drives), and a volume number (*volumenumber* is the volume ID of an available volume in the specified jukebox; it may be helpful to use the **vollist** command [refer to procedure **Using vollist to Display Volume Data** (subsequent section of this lesson)] to identify a suitable volume, such as a volume in the Space Pool, to use for this test).
- AMASS returns information on the health of library components in the following format:

```
--- CHECKING LIBRARY COMPONENTS ---
```

```
-api has been opened properly and AMASS is running.  
-mapping shared memory.  
-verifying the juke number.  
-validating volume number.  
-validating drive number and checking for active drive/s/.  
-saving the volume's status before inactivating it.  
-proceeding with physical test.  
-restoring signals and exiting.
```

```
--- CHECK COMPLETED!! ---
```

- If an error message is returned, it may be necessary to restart AMASS (refer to procedure **Rebooting AMASS** [previous section of this lesson]) and/or to check for possible hardware problems with drives or other components.

A command provided to display the status of the AMASS I/O activity is **sysperf**. This command returns several items:

- the number of reads and writes that are outstanding.
- the number of volumes (for reads) or volume groups (for writes) that are going to be used by those reads and writes.
- the current volumes in the drives.
- the I/O rate in Kb per second since the last update. This value first appears as a zero. Then AMASS continues to update the information at intervals based on a value for *updateinterval* entered by the operator.

Sysperf can often show the first sign of trouble. For example, if there are reads and writes in process but throughput is always 0, a problem is indicated. The most common problems are volumes and drives that go off line and/or inactive.

To run **sysperf**, use the following procedure.

Use *sysperf* to Display the Status of AMASS I/O Activity

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01).
- 2 Type **/usr/amass/bin/sysperf -k 5** and press the **Return/Enter** key.
 - The screen updates every 5 seconds and display information on the amass kernel (**-k**) in a form similar to the following example (**Note**: A different number of seconds may be entered to specify a different refresh rate.):

```
SYSTEM STATISTICS - Thu Sep 27 08:17:33
UPDATE INTERVAL   - 10 SEC
AVERAGE THROUGHPUT - 0 KBYTES/SEC

READ REQUESTS      # OF VOLUMES
      0              0

WRITE REQUESTS     # OF VOL GROUPS
      0              0

CACHE BLOCKS      2957 Total      2957 Free      0 Dirty
FNODES            800 Total      796 Free      4 Used

JUKE  DRIVE  VOLFLAGS  VOLUME  VOLGRP  KBYTES/SEC
```

- 3 To break out of the command, use **ctrl-c** (while holding down the **Control Key**, press **c**).
 - The screen stops updating and displays a UNIX prompt.

Volumes are monitored using the **vollist** command. For example, to show data on a particular volume (e.g., volume 100) use the following procedure.

Use *vollist* to Display Volume Data

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01).
 - 2 Type **/usr/amass/bin/vollist 100** (for this example, to specify volume 100).
 - AMASS displays the following:
- | VOL
NUM | VOL
GRP | JUKE
NUM | POS | VOL
LABEL | FLAGS | USED
(MB) | AVAIL
(MB) | DEAD
(%) | ERRS |
|------------|------------|-------------|-----|--------------|-------|--------------|---------------|-------------|------|
| 100 | 500 | 3 | NET | SD0060 | O | 99213 | 3167 | 0 | 0 |
- Note: In this example, the **O** in the **FLAGS** column indicates that the volume is offline.
- 3 To put volume 100 back on line, type **/usr/amass/bin/volloc -n 100**.

4 Type **/usr/amass/bin/vollist 100**.

- AMASS displays the following:

VOL NUM	VOL GRP	JUKE NUM	POS NET	VOL LABEL	FLAGS	USED (MB)	AVAIL (MB)	DEAD (%)	ERRS
100	500	3	NET	SD0060	A	99213	3167	0	0

- Note: In this example, the **A** in the **FLAGS** column indicates that the volume is now on line and Active.

If the output of **vollist** indicates that the volume is inactivated (i.e., there is an **I** in the **FLAGS** column), use the **amass_log** script to determine the nature of the problem. The **amass_log** script displays AMASS messages from the system log file. This script can provide helpful information under several circumstances, such as when a command gives unexpected results or when AMASS appears not to be functioning properly in other ways. Use the following procedure to run **amass_log**.

Use the *amass_log* script to Display AMASS Errors

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01).
 - 2 To change to the AMASS tools directory, type **cd /usr/amass/tools**, and then press the **Return/Enter** key.
 - The working directory is changed to **/usr/amass/tools**.
 - 3 Type **./amass_log logfilepath**, where **logfilepath** is the full pathname of the system log file to scan for AMASS messages, and then press the **Return/Enter** key.
 - On a Sun, the **logfilepath** is likely to be **/var/adm/messages**; on an SGI, the **logfilepath** is likely to be **/var/adm/SYSLOG**. Any AMASS error messages in the scanned log file are displayed.
 - 4 Perform the action recommended for the error message in the log.
 - The *AMASS System Administrator's Guide* (available electronically on **drg** servers [e.g., g0drg01, e0drg11, l0drg01, n0drg01] in directory **/usr/amass/books**) provides detailed information concerning error messages. An error message informs of critical problems that prevent AMASS from functioning. An error message is usually followed by a correction message, which provides instructions for correcting the situation. Sometimes, there is a previous warning message that may provide an accompanying correction message. Other messages may be identified by number only; the *System Administrator's Guide* provides a reference list, with accompanying corrective actions.
-

Unless use of the **amass_log** script shows that there are many errors on a volume that has been inactivated, you can reactivate the volume using the command:

/usr/amass/bin/volstat -a 100

(for this example, to reactivate volume 100).

Just as **vollist** provides information on the status of volumes, the command **drivelist** displays the status of drives available to AMASS. Active drives are noted by an **A**, and inactive drives are noted by an **I**. The command is **/usr/amass/bin/drivelist**. If AMASS inactivates a drive, use the **amass_log** script as described previously. Unless there is a hardware problem and several attempts have been made to ready the drive, it is usually appropriate to reactivate the drive using the **drivestat** command. For example, to reactivate drive 1 in jukebox 1, type the command **/usr/amass/bin/drivestat -a 1 1**.

A useful library utility included with AMASS is **quedisplay**. This utility permits the operator to see what is in the queue, and to diagnose problems such as the following:

- During an attempt to write to a file, the drive light does not illuminate.
- The system is slowing down.
- An AMASS command does not complete.

Figure 28 shows an example of the form of the output of the **quedisplay** utility. The output shows the queue, which consists of read and write requests, AMASS administration commands, and a list of libraries, drives, and what volumes they manage.

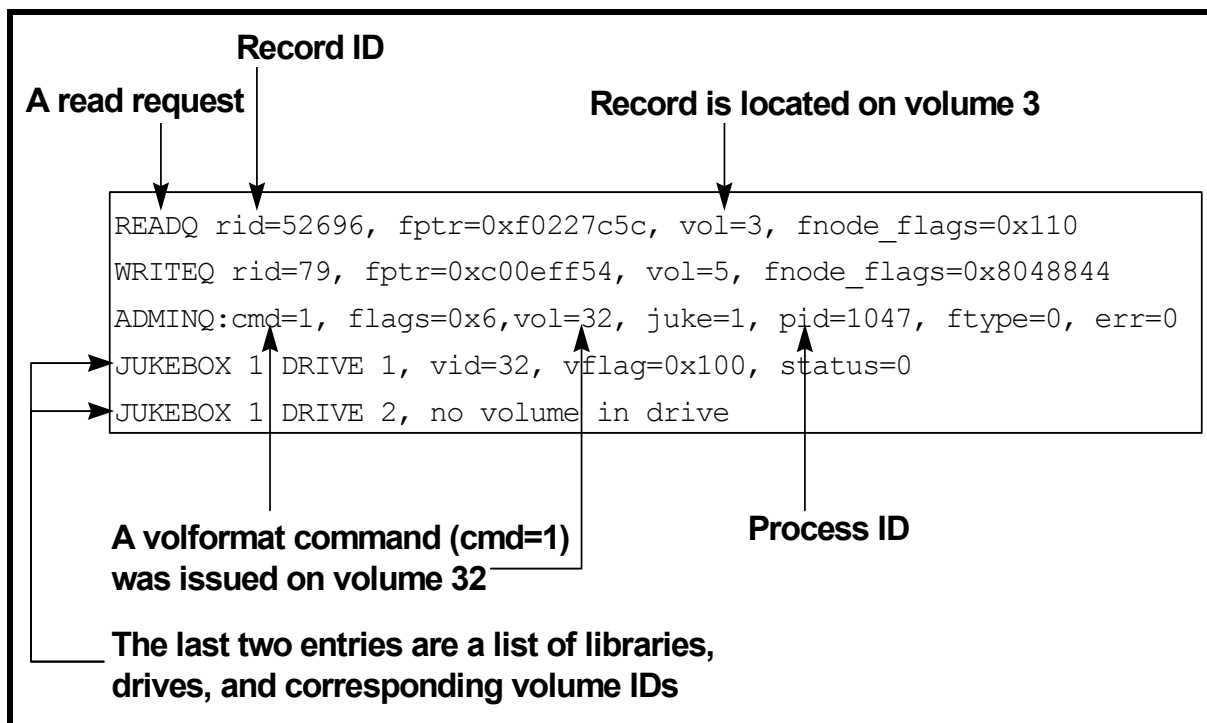


Figure 28. Sample output from AMASS quedisplay utility

Use the following procedure to monitor what is in the queue.

Use *quedisplay* to View What is in the AMASS Queue

- 1 Log in to the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01).
- 2 To change to the utilities directory, type **cd /usr/amass/utls**, and then press the **Return/Enter** key.
 - The working directory is changed to **/usr/amass/utls**.
- 3 Type **quedisplay**, and then press the **Return/Enter** key.
 - The AMASS queue is displayed in the following form:

```
READQ rid=52696, fptr=0xf0227c5c, vol=3, fnode_flags=0x110
WRITEQ rid=79, fptr=0xc00eff54, vol=5, fnode_flags=0x8048844
ADMINQ:cmd=1, flags=0x6,vol=32, juke=1, pid=1047, ftype=0, err=0
JUKEBOX 1 DRIVE 1, vid=32, vflag=0x100, status=0
JUKEBOX 1 DRIVE 2, no volume in drive
```

Note: In the output, "rid" = Record ID, "pid" = Process ID

If there are **READQ** or **WRITEQ** entries, the name(s) of the file(s) being processed can be determined by using the **filepath** command and the first number in the entry. For example, type **/usr/amass/utls/filepath 52696** for the first file number in the sample output shown previously in Figure 28.

Occasionally, a robot may lose synchrony with AMASS as to the location of media. The best way to verify this is to compare **quedisplay** and **medialist**. The **medialist** utility is a standalone program that communicates with the robot controller in the Powderhorn to determine the robot's view of media and their slot locations. If the two programs disagree, you can bring the two programs into synchrony using **mediamove**. Use the following procedure.

Use *mediamove* to Establish Synchrony Between *quedisplay* and *medialist*

- 1 Log in as **amass** at the FSMS host (e0drg11, g0drg01, l0drg01, or n0drg01).
- 2 Type **/usr/amass/utls/quedisplay** and then press the **Return/Enter** key.
 - AMASS displays the following information (for example of incorrect status).

```
. . .
JUKEBOX 1 DRIVE 1, no volume in drive
JUKEBOX 1 DRIVE 2, vid=50, vflags=0x4, status=0
```
- 3 Type **/usr/amass/utls/medialist** and then press the **Return/Enter** key.
 - AMASS displays the following information (for example of actual status).

```
. . .
SLOT VSD0098 FULL
DRIVE 1 FULL FROM VSD0096
DRIVE 2 FULL FROM VSD0097
```

- Note that the **medialist** result shows that drive 1 actually is occupied, although **quedisplay** registers that drive 1 is empty.
- 4 Type **/usr/amass/utills/mediamove 1 VSD0096 1** and then press the **Return/Enter** key.
- AMASS moves the volume from the *source* (drive **1** in this example) to the *destination* (slot **VSD0096** in this example) in the specified *jukeboxnumber* (jukebox **1** in this example), thereby bringing the actual status of drive 1 (as known by *medialist*) to the status reflected by *quedisplay*.
-

Recovery from Failure to Store or Retrieve Data

Successful data storage and retrieval functions are the heart of ECS. Successful ingest of data or processing of data to produce new science data granules require that Storage Management (STMGT) is inserting the product into the archive and that Science Data Server (SDSRV) is inserting the associated metadata into the inventory. Staging disks and cache managers for the Archive server and the FTP server are also involved in this process. To check the functioning of these elements, it is necessary that the ESDTs for the data to be inserted are installed and available, and that subscriptions have been registered.

Troubleshooting failures to store or retrieve data (as well as other failures) often requires review of server or application log files. This section contains a general procedure for reviewing log files to check for proper start-up and communications. It also has a procedure for a special case of reviewing log files for the Storage Management Request Manager server, and a procedure for reviewing the current **tac** log file of interactions between AMASS and ACSLS. Separate procedures then address recovery from failure to insert (store) data and recovery from failure to acquire (retrieve) data.

Checking Server Log Files

Use the following procedure for checking server log files for nominal start-up and communications.

Checking Server Log Files

- 1 Log in to the host for the server and log(s) to be examined.
- 2 Type **cd /usr/ecs/<MODE>/CUSTOM/logs** and then press the **Return/Enter** key.
 - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/logs**.
- 3 To view a server log, type **pg filename** and then press the **Return/Enter** key.
 - *filename* refers to the log file to be reviewed (e.g., **EcDsScienceDataServer.ALOG**, **EcDsScienceDataServerDebug.log**).
 - The first page of the log file is displayed; additional sequential pages can be displayed by pressing the **Return/Enter** key at the **:** prompt.

- Although this procedure has been written for the **pg** command, any UNIX editor or visualizing command (e.g., **vi**, **view**, **more**, **tail**) can be used to review the log file.
- Typically, the **<server>Debug.log** captures more detailed information than the **<server>.ALOG**. However, for some servers (e.g., **SDSRV**), there may be significant detail in the **<server>.ALOG**. It is also important to note that the **DebugLevel** parameter setting in the **Configuration Registry** determines the level of detail captured in the **<server>Debug.log** (**0** is off, a setting of **1** captures status and errors, a setting of **2** captures major events, and a setting of **3** is a full trace recording of all activity). If the **DebugLevel** has been set to one of the lower levels during operations, the System Administrator may set it to **3** during troubleshooting. Similarly, the **AppLogLevel** parameter setting determines the level of detail captured in the **<server>.ALOG** (**0** provides a full trace recording of all events, **1** provides messages related to all major events, **2** yields just records of errors, and **3** turns recording off). (Note: There are other debug levels available for some logs; Storage Management (STMGT) offers "enhanced" debugging based on bitmasks. Level **7** provides a four-bit level for detailed database debugging. Level **15** provides an eight-bit level that repeatedly dumps the in-memory request queue in the STMGT Request Manager.

4 Review the log file(s) to determine if there are any indications of connection problems or errors at start up.

- The log file for the called server may contain an error message indicating a problem at start-up. The debug log should indicate a typical start sequence, including (sample log entries in the following material were taken from a debug log showing start-up for **EcDsStFtpServer**):

- Get parameters from registry (log entries similar to the following).

```
DSS EcDsStFtpServer Server Debug log on f2acg01 starting at Mon
Jun  4 07:57:45 EDT 2001
EcAgInstanceID Sequence Number is 3870
Setting up environment variables needed for DCE:
RPC_UNSUPPORTED_NETIFS = ""
/usr/ecs/DEV07/CUSTOM/bin/DSS/EcDsStFtpServer ConfigFile
/usr/ecs/DEV07/CUSTOM/cfg/EcDsStFtpServer.CFG ecs_mode DEV07
StartTemperature cold
Started process EcDsStFtpServer in mode DEV07 with PID 2709893
EcRgRegistry_1_0::ctor this = 0x104eef38
EcRgRegistry_1_0::ctor this = 0x104eef88
FoIpPtToPtPortalImp::Send sent 20/20
FoIpPtToPtPortalImp::Send sent 219/219
FoIpPtToPtPortalImp::Receive got 20
FoIpPtToPtPortalImp::Receive got 1024
FoIpPtToPtPortalImp::Receive got 246
***** After Retrieving of RGY: Name = EcDsStFtpServerNONE
ProgramID = 4645102
ApplicationID = 4600000
Release = B
DeltaTime = 0
Site = RBD
```



```

SubSysName = DSS
MajorVersion = 1
MinorVersion = 0
DebugLevel = 3
AppLogLevel = 0
AppLogSize = 3000000
DBServer = f2acg01_srvr
DBLoginName = EcDsStFtpServer
DBName = stmgtdbl

```

- Load resource catalogs (log entries indicate the loading, or that the loading did not complete, similar to the following).

```

06/04/01 07:57:47: Thread ID : 65536 : loading resource catalog
file from
/usr/ecs/DEV07/CUSTOM/data/DSS/ResourceCatalogs/DsMdResource.dat.r
cat
06/04/01 07:57:48: Thread ID : 65536 : loading resource catalog
file from
/usr/ecs/DEV07/CUSTOM/data/DSS/ResourceCatalogs/EcDsSdHr.dat.rcat
06/04/01 07:57:48: Thread ID : 65536 : loading resource catalog
file from
/usr/ecs/DEV07/CUSTOM/data/DSS/ResourceCatalogs/DsSrResource.dat.r
cat
06/04/01 07:57:48: Thread ID : 65536 : loading resource catalog
file from
/usr/ecs/DEV07/CUSTOM/data/DSS/ResourceCatalogs/DsGlResource.dat.r
cat
06/04/01 07:57:48: Thread ID : 65536 : loading resource catalog
file from
/usr/ecs/DEV07/CUSTOM/data/DSS/ResourceCatalogs/DsShResource.dat.r
cat
06/04/01 07:57:48: Thread ID : 65536 : loading resource catalog
file from
/usr/ecs/DEV07/CUSTOM/data/DSS/ResourceCatalogs/EcDsSdHc.dat.rcat

```

- Pre-cache errors associated with database connectivity (log entries similar to the following).

```

06/04/01 07:57:48: Thread ID : 65536 : User Name      :
EcDsStFtpServer | Thread 65536
06/04/01 07:57:48: Thread ID : 65536 : Database Name      :
stmgtdbl_DEV07 | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : Server Name      :
f2acg01_srvr | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : DsShTSStorage: creating
the MutexVec for this thread
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30141 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING: DsEstUnknownError
(30141) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30143 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING: DsEstUnknownError
(30143) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30139 (Not
found) | Thread 65536

```

```

06/04/01 07:57:49: Thread ID : 65536 : CACHING: DsEstUnknownError
(30139) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30142 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING: DsEstUnknownError
(30142) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30148 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING: DsEstUnknownError
(30148) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30144 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING: DsEstUnknownError
(30144) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30145 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING: DsEstUnknownError
(30145) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30147 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING: DsEstUnknownError
(30147) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30146 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING: DsEstUnknownError
(30146) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30211 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING: DsEstUnknownError
(30211) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : SEARCHING FOR: 30140 (Not
found) | Thread 65536
06/04/01 07:57:49: Thread ID : 65536 : CACHING: DsEstUnknownError
(30140) | Thread 65536

```

- Get server configuration parameters from the database (log entries similar to the following).

```

06/04/01 07:57:49: Thread ID : 65536 : BaseReal::Ctor: Server
Name is - EcDsStFtpServerNONE | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : User Name      :
EcDsStFtpServer | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : Database Name      :
stmgtdb1_DEV07 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : Server Name      :
f2acg01_srvr | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : myTransactionList[0]: use
stmgtdb1_DEV07 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: use stmgtdb1_DEV07 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : myTransactionList[1]: exec
DsStCPSelectByName "EcDsStFtpServerNONE" | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: exec DsStCPSelectByName "EcDsStFtpServerNONE" | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 :
DBIF:Fetch:[8.000000][EcDsStFtpServerNONE][1][10][FTP][ ][0][FTPA
][NONE][4194304] | Thread 65536

```

- Spawn receptionist thread and register server in the database (log entries similar to the following).

```
06/04/01 07:57:50: Thread ID : 65536 :
DsStReceptionist:BindSocketGetInfo: Port assigned is 13441 |
Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : myTransactionList[0]: use
stmgtdbl_DEV07 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: use stmgtdbl_DEV07 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : myTransactionList[1]: exec
DsStCPRegisterServer 8, 13441, "f2acg01" | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: exec DsStCPRegisterServer 8, 13441, "f2acg01" | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : Ftp:Ctor:
EcDsStFtpServerNONE | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : myTransactionList[0]: use
stmgtdbl_DEV07 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: use stmgtdbl_DEV07 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : myTransactionList[1]: exec
DsStCPSelectById 8 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: exec DsStCPSelectById 8 | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 :
DBIF:Fetched:[8.000000][EcDsStFtpServerNONE][1][10][FTP][f2acg01][
13441][FTP][NONE][4194304] | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : Ftp:Ctor: Leaving | Thread
65536
```

- Spawn service threads (log entries similar to the following).

```
06/04/01 07:57:50: Thread ID : 65536 : Ftp:Startup: temperature =
cold | Thread 65536
06/04/01 07:57:50: Performing startup processing | Thread 65536
06/04/01 07:57:50: Thread ID : 65536 : Spawning service threads |
Thread 65536
06/04/01 07:57:50: Thread ID : 65536 :
BR:GetThreadPoolConfiguration | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : myTransactionList[0]: use
stmgtdbl_DEV07 | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: use stmgtdbl_DEV07 | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : myTransactionList[1]: exec
DsStSTCSelectForServer 8, "ThreadPool" | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: exec DsStSTCSelectForServer 8, "ThreadPool" | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : DBIF:Fetched:[ThreadPool
][10][0][0][0][0][10] | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : 3_2709893_0757-
1125858625_155062001_f2acg01:FTP: BR:GetThreadPoolConfiguration
Returning | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : Ftp: Spawning a service
thread | Thread 65536
06/04/01 07:57:51: Starting a new service thread | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : Ftp: Spawning a service
thread | Thread 65536
```

```

06/04/01 07:57:51: 06/04/01 07:57:51: Thread ID : 65554 : Waiting
for work | Thread 65554
06/04/01 07:57:51: Thread ID : 65554 : DsShTSSStorage: creating
the MutexVec for this thread
06/04/01 07:57:51: Thread ID : 65554 : Waking up manager thread |
Thread 65554
.
.
06/04/01 07:57:51: Starting a new service thread | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : Ftp: Spawning a service
thread | Thread 65536
06/04/01 07:57:51: Starting a new service thread | Thread 65536
06/04/01 07:57:51: Thread ID : 65559 : Waiting for work | Thread
65559
06/04/01 07:57:51: Thread ID : 65560 : Waiting for work | Thread
65560
06/04/01 07:57:51: Thread ID : 65561 : Waiting for work | Thread
65561

```

- **Process Restart Notification for server restart ("Ready to accept requests") (log entries similar to the following).**

```

06/04/01 07:57:51: Thread ID : 65536 : myTransactionList[0]: use
stmgtdbl_DEV07 | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: use stmgtdbl_DEV07 | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : myTransactionList[1]:
BEGIN TRANSACTION OUTER_278888352 | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: BEGIN TRANSACTION OUTER_278888352 | Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : myTransactionList[2]: exec
DsStGRRestartNotification "10_2709893_0757-1125858625_15506
2001_f2acg01:FTPA:Server restart", "EcDsStFtpServerNONE", "cold" |
Thread 65536
06/04/01 07:57:51: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: exec DsStGRRestartNotification "10_2709893_0757-1125858625
_155062001_f2acg01:FTPA:Server restart", "EcDsStFtpServerNONE",
"cold" | Thread 65536
06/04/01 07:57:52: Thread ID : 65536 : DBIF:Fetched:[] | Thread
65536
06/04/01 07:57:52: Thread ID : 65536 : DBIF:Fetched:[8.000000] |
Thread 65536
06/04/01 07:57:52: Thread ID : 65536 : DBIF:Execute: Ultimate
SQL: COMMIT TRANSACTION OUTER_278888352 | Thread 65536
06/04/01 07:57:52: Thread ID : 0 : No servers to awaken -- get
status | Thread 0
06/04/01 07:57:52: Thread ID : 65536 : Spawning manager thread |
Thread 65536
06/04/01 07:57:52: Ready to accept requests | Thread 65564

```

- **Check queue for requests ("Waiting for an event" means there is nothing else in the queue.) (log entries similar to the following).**

```

06/04/01 07:57:52: Thread ID : 65564 :
BR:ProcessCancelledRequests | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : DsShTSSStorage: creating
the MutexVec for this thread

```

```

06/04/01 07:57:52: Thread ID : 65564 : myTransactionList[0]: use
stmgtdbl_DEV07 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : DBIF:Execute: Ultimate
SQL: use stmgtdbl_DEV07 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : myTransactionList[1]: exec
DsStGRSelectCancelled 8 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : DBIF:Execute: Ultimate
SQL: exec DsStGRSelectCancelled 8 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 :
BR:ProcessCancelledRequests: Nothing cancelled | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 :
BR:ProcessCancelledRequests Returning | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : Ftp: Getting next request
| Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : myTransactionList[0]: use
stmgtdbl_DEV07 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : DBIF:Execute: Ultimate
SQL: use stmgtdbl_DEV07 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : myTransactionList[1]: exec
DsStFRGetNextRequest 8 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : DBIF:Execute: Ultimate
SQL: exec DsStFRGetNextRequest 8 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : GetNextRequest: No
requests available | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : Waiting for an event |
Thread 65564

```

- The log file for the server from which the call originated may indicate a problem completing a connection. The log should indicate successful awakening of a remote host, with entries similar to the following:

```

06/04/01 07:57:52: Thread ID : 65536 :
DsStPatron:AwakenRemoteServer: Hostname - f2acg01 | Thread 65536
06/04/01 07:57:52: Thread ID : 65536 :
DsStPatron:AwakenRemoteServer: Port Number - 13441 | Thread 65536
06/04/01 07:57:52: Thread ID : 65536 : Patron: Creating new entry
for EcDsStFtpServerNONE | Thread 65536
06/04/01 07:57:52: Thread ID : 65536 : Trying gethostbyname_r() 0
of 5 attempts | Thread 65536
06/04/01 07:57:52: Thread ID : 65536 : Waking up
EcDsStFtpServerNONE | Thread 65536

```

and should indicate completion of a connection to the called server, with entries similar to the following:

```

06/04/01 07:57:52: Thread ID : 65553 :
DsStReceptionist:WaitForConnections: A connection has been
accepted | Thread 65553
06/04/01 07:57:52: Thread ID : 65564 :
BR:ProcessCancelledRequests | Thread 65564
06/04/01 07:57:52:
06/04/01 07:57:52: Thread ID : 65553 : Waking up manager thread |
Thread 65553
06/04/01 07:57:52: Thread ID : 65564 : : 06/04/01 07:57:52: read
ID : 7:57:52: DsShTSStorageDsShTSStorage: creating the MutexVec
for this thread: creating the MutexVec for this thread665553: 53 :
DsShTSStorage: creating the MutexVec for this thread

```

```

06/04/01 07:57:52: Thread ID : 65564 : myTransactionList[0]: use
stmgtdbl_DEV07 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : DBIF:Execute: Ultimate
SQL: use stmgtdbl_DEV07 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : myTransactionList[1]: exec
DsStGRSelectCancelled 8 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 : DBIF:Execute: Ultimate
SQL: exec DsStGRSelectCancelled 8 | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 :
BR:ProcessCancelledRequests: Nothing cancelled | Thread 65564
06/04/01 07:57:52: Thread ID : 65564 :
BR:ProcessCancelledRequests Returning | Thread 65564.

```

- This procedure is applicable for reviewing logs for different types of errors and events on ECS servers.

5 Exit the log file (e.g., from **pg**, type **q** and then press the **Return/Enter** key).

A Special Case: Checking the Request Manager Server Debug Log

The Request Manager server in the Storage Management computer software configuration item of the Data Server Subsystem processes requests from external clients (processes outside of Storage Management). Requests between Storage Management servers are passed directly from one server to another.

- Requests that require one of the Storage Management servers to perform processing are checkpointed (except requests that can be serviced solely through SQL).
 - Checkpointing involves recording the request's state (e.g., "checkpointed," "failed," "completed") in the database to assist in error recovery.
- Requests that can be serviced solely through SQL are considered "trivial" requests.
 - Trivial requests are not checkpointed.
 - Examples include attaching to a staging disk, getting capacity, and getting block size.
 - Trivial requests submitted from outside Storage Management are serviced by the Request Manager server.
 - Trivial requests originating within Storage Management are passed directly from the client to the database server.

The Request Manager server (like other Storage Management servers) can manage several concurrent activities. This is accomplished through the use of threads. There are several different kinds of threads:

- Manager thread.
 - One per Storage Management server.

- Responsible for dequeuing requests and assigning them to service threads.
- Checks for cancelled requests.
- Service thread.
 - Multiple threads per Storage Management server.
 - Responsible for the actual servicing of requests.
 - Logs all progress including all changes of request state.
 - Notifies submitter when request has been completed.
- Receptionist thread.
 - One per Storage Management server.
 - Registers the server as "up" in the database.
 - Sits on a socket, waiting for connections from other Storage Management servers.
 - Unregisters the server at shutdown.
- Inbound RPC thread.
 - Spawned by a request from a Storage Management client.
 - Hands off the request to the manager thread and waits for completion of the request.
- Housekeeper thread.
 - Watches for completed requests that have not previously been seen and processed.

Information concerning Request Manager server processing of requests (identified by thread) is recorded in the Request Manager server debug log (assuming some level of debug log recording is specified in the Registry database).

Trivial requests typically involve the following types of activities:

- Inbound RPC thread appears with a request.
- Manager thread dequeues the request and assigns it to a service thread.
- Service thread recognizes the thread as "trivial."
 - A "No checkpointing required -- going straight to responded" message is recorded in the Request Manager server debug log.
- Service thread executes the database transaction for results.
 - When the request is completed, a "Done servicing" message is recorded in the Request Manager server debug log.

- If the request fails, an "Unable to service" message is recorded in the Request Manager server debug log.
- Service thread hands the results to the inbound RPC thread.
 - A "Notifying the client" message is recorded in the Request Manager server debug log.
- Inbound RPC thread silently returns to the client with the results.

Non-trivial requests are forwarded to the appropriate Storage Management server (e.g., EcDsStFtpServer, EcDsStStagingDiskServer, EcDsStArchiveServer) for processing.

- Some of the same types of entries are made in the Request Manager server debug log for non-trivial requests as for trivial requests.
 - For example:
 - "Waking up service thread" (Request Manager is preparing to process the request).
 - "Done servicing" (request processing has been completed).
 - "Unable to service" (the request has failed).
- Although some trivial requests include "token" statements, tokens are characteristic of non-trivial requests.
 - A token includes request information that varies with the type of operation to be performed.
 - For example, a token for an ftp request might include the following types of data:
 - Stored procedure (e.g., DsStFRInsert) [other types of stored procedures include DsStSDRInsert and DsStGRMapLogicalArchiveId].
 - RPC ID (e.g., RPCId=1821_535_1109-1124464729_171062001_x0ins01.xdc.ecs.nasa.gov:SBSVSDSV1DSDD1DSDD4:).
 - Username.
 - Encrypted password.
 - Host.
 - Source path.
 - Destination path.
 - External request ID.
 - Server name (e.g., EcDsStFtpServerNONE) [other types of operations might involve the EcDsStStagingDiskServerDRP1 for example].

- Type of operation (e.g., FtpPush) [other types of operations include ArRetrieve, SDAllocateDisk, SDLinkFile].
 - Submitter (e.g., DSDD) [other types of operations might involve SDSV].
 - Priority.
- The server to which the request was sent is identified by name (ServerName).
 - Transaction ID is embedded in the RPC ID (the portion before the first colon in the RPC ID).

A "transaction" may involve multiple operations on a host or several hosts. Consequently, multiple threads may be used on each relevant host. Use the following procedure to check the debug log for the Storage Management Request Manager server.

Checking the Request Manager Server Debug Log

- 1** Log in to the Distribution Server host (e.g., e0dis02, g0dis02, l0dis02, n0dis02).
- 2** To change to the logs directory, type **cd /usr/ecs/<MODE>/CUSTOM/logs** then press the **Return/Enter** key.
 - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/logs**.
- 3** Type **pg filename** then press the **Return/Enter** key.
 - **filename** refers to the appropriate Request Manager debug log.
 - For example: **pg EcDsStRequestManagerServerDebug.log**
 - The content of the first page of the specified file is displayed.
 - Although this procedure has been written for the **pg** command, any UNIX editor or visualizing command (e.g., **vi**, **view**, **more**) can be used to review the log file.
- 4** At the **:** prompt type **/date time** then press the **Return/Enter** key.
 - **date time** refers to the approximate date and time of the problem.
 - For example:
/06/18/01 12:17:31
 - The file is searched for the specified text.
 - If the specified text is in the log file, the following type of response is displayed.

```
...skipping forward
06/18/01 12:17:31: Thread ID : 105 : DsShTSStorage: creating the
MutexVec for this thread
[...]
```
 - If the specified text is not in the log file, the following type of response is displayed.

Pattern not found:

- If the specified text is not in the log file, verify the following aspects of Steps 3 and 4:
 - Date and time were entered correctly (Step 4).
 - Proper file was opened (Step 3).

5 At the `:` prompt type **/Unable to service** then press the **Return/Enter** key.

- **pg** searches the file for the specified text.
 - If the specified text is in the log file, the following type of response is displayed.

```
...skipping forward
2:IngestRQ409GR1 Unable to service | Thread 52
[...]
```

- If the specified text is not in the log file, the following type of response is displayed.

Pattern not found:

- If the specified text is in the file, go to Step 7.
- If the specified text is not in the file, go to Step 6.

6 Examine the contents of the log file to determine which thread is associated with the problem being investigated.

- The following **pg** commands (at the `:` prompt) are useful:
 - ***n*** then **Return/Enter** (go to Page *n*).
 - **Return/Enter** or **+1** then **Return/Enter** (go down to the next page).
 - **-1** then **Return/Enter** (go back to the preceding page).
 - **+*n*** then **Return/Enter** (go down *n* number of pages).
 - **-*n*** then **Return/Enter** (go back *n* number of pages).
 - **+*nl*** then **Return/Enter** (go down *n* number of lines).
 - **-*nl*** then **Return/Enter** (go back *n* number of lines).
 - **q** then **Return/Enter** (exit from **pg**).

7 At the `:` prompt type the appropriate text (depending on the direction of the desired search) then press the **Return/Enter** key:

- 8 To search back toward the beginning of the file, type **^Waking up service thread *n*^** and then press **Return/Enter**.
- 9 To search toward the end of the file, type **/Waking up service thread *n*** and then press **Return/Enter**.

- For example:

^Waking up service thread 52^

- The file is searched back toward the beginning of the file for the specified text.

- If the specified text is in the log file, the following type of response is displayed.

```
...skipping backward
06/18/01 12:17:31: Thread ID : 102 : Waking up service thread 52
| Thread 102
[...]
```

- If the specified text is not in the log file, the following type of response is displayed.

```
Pattern not found:
```

- The entries "Waking up service thread *n*" and "Unable to service | Thread *n*" bracket the thread servicing in which an error occurred.

NOTE: Thread IDs are reused frequently. There are likely to be many processes with the same thread ID in any particular log file. It is important to follow the correct instance of the thread.

NOTE: It is likely that the Request Manager would try again to process a failed request. Subsequent request processing may use the same thread ID or a different thread ID. However, it would involve the same transaction ID.

- A "No checkpointing required -- going straight to responded" entry associated with the thread ID indicates that the request is "trivial."

10 At the : prompt type **/SEARCHING** then press **Return/Enter**.

- The file is searched for the specified text.
 - If the specified text is in the log file, the following type of response is displayed.

```
...skipping forward
06/18/01 12:17:31: Thread ID : 52 : SEARCHING FOR: 30148 (Found)
| Thread 52
06/18/01 12:17:31: Thread ID : 52 : SEARCHING FOR: 30148 (Found)
| Thread 52
06/18/01 12:17:31: Thread ID : 52 : DsStStoredProcedures::Execute
- ERROR: Could not execute stored procedure | Thread 52
06/18/01 12:17:31: Thread ID : 52 : Error encountered in stored
procedure | Thread 52
06/18/01 12:17:31: Thread ID : 52 : DBIF:Execute: Ultimate SQL:
ROLLBACK TRANSACTION OUTER_7077776 | Thread 52
06/18/01 12:17:32: Thread ID : 52 : 1_4501810_1217-
1124633447_169062001_p0icg01.pvc.ecs.nasa.gov:IPOBIPOB1INRM1IGSA15
:IngestRQ409GR1 Done servicing | Thread 52
06/18/01 12:17:32: Thread ID : 52 : 1_4501810_1217-
1124633447_169062001_p0icg01.pvc.ecs.nasa.gov:IPOBIPOB1INRM1IGSA15
:IngestRQ409GR1 Unable to service | Thread 52
06/18/01 12:17:32: Thread ID : 52 : 1_4501810_1217-
1124633447_169062001_p0icg01.pvc.ecs.nasa.gov:IPOBIPOB1INRM1IGSA15
:IngestRQ409GR1 Marked as unassigned | Thread 52
06/18/01 12:17:32: Thread ID : 52 : 1_4501810_1217-
1124633447_169062001_p0icg01.pvc.ecs.nasa.gov:IPOBIPOB1INRM1IGSA15
:IngestRQ409GR1 Notifying the client | Thread 52
06/18/01 12:17:32: Thread ID : 52 : Waiting for work | Thread 52
06/18/01 12:17:32: Thread ID : 52 : Waking up manager thread |
Thread 52
[...]
```

- In the preceding example the expression **SEARCHING** is associated with Thread ID 52.
- The context of the **SEARCHING** statement indicates the type and source of the problem; in this case there appears to be a problem executing a stored procedure.
- If the specified text is not in the log file, the following type of response is displayed.

```
Pattern not found:
```

11 If the expression **SEARCHING** is not associated with the specified thread in the lines displayed, repeat Step 8.

- 12 If necessary, at the **:** prompt type **-2l** [lower-case letter l] then press the **Return/Enter** key.
- **pg** simulates scrolling the screen backward two lines (or any other number of lines that is typed at the prompt).
 - The file is redisplayed to include the two lines that preceded the page previously displayed.
 - For example:


```
...skipping backward
06/18/01 12:17:31: Thread ID : 52 : DBIF:Execute: Ultimate SQL:
exec DsStSDAttachDisk
"/usr/ecs/TS2/CUSTOM/pdps/x0spg01/data/DpPrRm/x0spg01_disk",
"SDSV", 0 | Thread 52
06/18/01 12:17:31: Thread ID : 52 : SEARCHING FOR: 30148 (Found)
| Thread 52
06/18/01 12:17:31: Thread ID : 52 : SEARCHING FOR: 30148 (Found)
| Thread 52
06/18/01 12:17:31: Thread ID : 52 : DsStStoredProcedures::Execute
- ERROR: Could not execute stored procedure | Thread 52
06/18/01 12:17:31: Thread ID : 52 : Error encountered in stored
procedure | Thread 52
[...]
```
 - The additional lines preceding "SEARCHING FOR" in the example indicate that the stored procedure in which the error was encountered is DsStSDAttachDisk.
- 13 At the **:** prompt type **q** then press the **Return/Enter** key.
- **pg** exits from the Request Manager server debug log file.
- 14 If the request is a trivial request, go to Step 22.
- 15 If the request is a non-trivial request, open a separate UNIX window.
- The results of related operations on the server involved in performing copy or ftp functions for the transaction are going to be checked in a separate UNIX window.
- 16 In the new UNIX window log in to the appropriate server host (e.g., e0drg11, g0drg01, l0drg01, n0drg01) for the server involved in performing copy or ftp functions for the transaction.
- 17 At the shell prompt type **grep 'TransactionId' filename | grep 'LogProgress'** then press the **Return/Enter** key.
- For example:


```
grep 'af610628-' EcDsStArchiveServerDebug.log | grep 'LogProgress'
```
 - *filename* refers to the name of the log file for the process involved in performing copy or ftp functions for the transaction.

- **TransactionId** refers to the Transaction ID associated with the applicable request.
- In this example af610628-1dd1-11b2-a047-af3a589fd88e is the relevant Transaction ID.
 - However, usually it is not necessary to use the entire Transaction ID in the command; a representative sample (e.g., af610628- from the example) should be sufficient.
 - References to other Transaction IDs and entries that do not contain the string "LogProgress" are filtered out so references to the specified Transaction ID that contain the string "LogProgress" are the only log entries displayed.
 - The string "LogProgress" is a filter for references to stored procedure DsStGRLogProgress.
 - Progress is logged for copy and ftp input/output at each block.
 - The following type of response is displayed:

```
06/26/01 12:46:00: Thread ID : 65674 : myTransactionList[1]: exec
DsStGRLogProgress "af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2S
C:MOD03.001:55732", 0, 1, "files" | Thread 65674
06/26/01 12:46:00: Thread ID : 65674 : DBIF:Execute: Ultimate
SQL: exec DsStGRLogProgress "af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2S
C:MOD03.001:55732", 0, 1, "files" | Thread 65674
06/26/01 12:46:43: Thread ID : 65674 : : 06/26/01 12:46:43: read
ID : 2:46:43: myTransactionmyTransactionList[1]: exec
DsStGRLogProgress "af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2S
C:MOD03.001:55732", 60, 60, "MB"List[1]: exec DsStGRLogProgress
"af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2S
C:MOD03.001:55732", 60, 60, "MB"65714read 65674 : 74
06/26/01 12:46:43: Thread ID : 65674 : DBIF:Execute: Ultimate
SQL: exec DsStGRLogProgress "af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2S
C:MOD03.001:55732", 60, 60, "MB"0DBIF:Execute: Ultimate SQL: exec
DsStGRLogProgress "af610628-1dd1-11b2-a047-
af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2S
C:MOD03.001:55732", 60, 60, "MB"06/26/01 12:46:43: 6/26/01
12:46:43: | Thread : 65714read 65674 : 74
```

- If no progress is indicated, go to Step 22.
- 18 Click in the UNIX window for the Distribution Server host.
- 19 Type **grep 'TransactionId' filename | grep 'Done servicing'** then press **Return/Enter**.
- **filename** refers to the appropriate Request Manager debug log.

- For example:

grep 'af610628-' EcDsStRequestManagerServerDebug.log | grep 'Done servicing'

- If the operation has been completed, the following type of response is displayed:

```
06/26/01 12:46:00: Thread ID : 52 : af610628-1dd1-11b2-a047-af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2SC:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:44: Thread ID : 52 : af610628-1dd1-11b2-a047-af3a589fd88e:PDPSSDSV1DSDD1DSDD10DSDD1DSDD1:MoPGE02#sy14182000TS2SC:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:45: Thread ID : 52 : af610628-1dd1-11b2-a047-af3a589fd88e:PDPSSDSV1DSDD1DSDD2DSDD1DSDD3:MoPGE02#sy14182000TS2SC:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:47: Thread ID : 52 : af610628-1dd1-11b2-a047-af3a589fd88e:PDPSSDSV1DSDD1DSDD2DSDD1DSDD3:MoPGE02#sy14182000TS2SC:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:47: Thread ID : 52 : af610628-1dd1-11b2-a047-af3a589fd88e:PDPSSDSV1DSDD1DSDD2DSDD1DSDD7:MoPGE02#sy14182000TS2SC:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:50: Thread ID : 52 : af610628-1dd1-11b2-a047-af3a589fd88e:PDPSSDSV1DSDD1DSDD2DSDD1DSDD7:MoPGE02#sy14182000TS2SC:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:51: Thread ID : 52 : af610628-1dd1-11b2-a047-af3a589fd88e:PDPSSDSV1DSDD1DSDD4:MoPGE02#sy14182000TS2SC:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:56: Thread ID : 52 : af610628-1dd1-11b2-a047-af3a589fd88e:PDPSSDSV1DSDD1DSDD4:MoPGE02#sy14182000TS2SC:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:56: Thread ID : 52 : af610628-1dd1-11b2-a047-af3a589fd88e:PDPSSDSV1DSDD1DSDD8:MoPGE02#sy14182000TS2SC:MOD03.001:55732 Done servicing | Thread 52
06/26/01 12:46:59: Thread ID : 52 : af610628-1dd1-11b2-a047-af3a589fd88e:PDPSSDSV1DSDD1DSDD8:MoPGE02#sy14182000TS2SC:MOD03.001:55732 Done servicing | Thread 52
```

- The statement "Done servicing" shows that the operation has been completed; however, it provides no indication as to whether the operation succeeded or failed.
- If "Done servicing" is followed by "Unable to service," (as described in Step 19) the operation failed.
- If the operation has not been completed, no file entries are displayed (the UNIX prompt is displayed).
 - It may just be slow to complete.
- If the operation has been completed, go to Step 19.
- If the operation has not been completed, go to Step 20.

- 20 Type **grep 'TransactionId' filename | grep 'Unable to service'** then press the **Return/Enter** key.
- *filename* refers to the appropriate Request Manager debug log.
 - For example:
grep '2a7d4168-' EcDsStRequestManagerServerDebug.log | grep 'Unable to service'
 - If the request has failed, the following type of response is displayed:


```
06/26/01 12:56:22: Thread ID : 52 : 2a7d4168-1dd2-11b2-8c52-99d0f708dce5:PDPSSDSV1:MoPGE02#sy14182000TS2MOD02OBC Unable to service | Thread 52
06/26/01 12:56:22: Thread ID : 52 : 2a7d4168-1dd2-11b2-8c52-99d0f708dce5:PDPSSDSV4:MoPGE02#sy14182000TS2MOD02OBC Unable to service | Thread 52
```

 - If the operation has failed, return to Step 7.
 - If the operation has not failed, no file entries are displayed (the UNIX prompt is displayed).
- 21 If the operation has not failed, at the shell prompt type **tail -f filename | grep 'TransactionId'** and then press the **Return/Enter** key.
- *filename* refers to the appropriate Request Manager debug log.
 - *TransactionId* refers to the Transaction ID associated with the applicable request.
 - For example:
tail -f EcDsStRequestManagerServerDebug.log | grep 'af610628-'
 - If new entries are being posted to the log, the operation has not finished yet.
 - If the same entries continue to be repeated over and over, it may be necessary to restart the server.
 - If it is necessary to exit from a tailed log, type **ctrl-c** (while holding down the **Control Key**, press **c**).
- 22 If the operation has not finished yet, monitor the tailed log for awhile.
- If the operation does not seem to finish (i.e., if entries continue to be made to the tailed log) after a reasonable period of time (e.g., 30 minutes), it may be necessary to restart the Request Manager server.
 - If it is necessary to exit from a tailed log, type **ctrl-c** (while holding down the **Control Key**, press **c**).

- 23 If problems were detected in the Request Manager server debug log and/or the log file for the process involved in performing copy or ftp functions for the transaction, it may be necessary to restart the server(s) performing those functions.
- If server restart does not resolve the problem, it is appropriate to notify the Help Desk and prepare a Trouble Ticket.
- 24 If no problems were detected in the Request Manager server debug log or the log file for the process involved in performing copy or ftp functions for the transaction, check the Science Data Server log files; use Procedure 17.7.2.1 for **Checking Server Log Files**.
-

Checking the *tac* Log

Each day a current **tac_00** log on the FSMS host records interactions between AMASS and ACSLS. This log can provide helpful information in troubleshooting problems manifested in those interactions. Use the following procedure to check the **tac_00** log.

Checking the *tac* Log

- 1 Log in as **amass** at the FSMS host.
- 2 Type **cd /usr/amass/logs/tac** and then press the **Return/Enter** key.
 - The working directory is changed to **/usr/amass/logs/tac**.
- 3 Use the current tac log to investigate possible problems in communication between AMASS and ACSLS. To view the current tac log, type **pg tac_00** and then press the **Return/Enter** key.
 - The first page of the log file is displayed; additional sequential pages can be displayed by pressing the **Return/Enter** key at the **:** prompt.
 - Although this procedure has been written for the **pg** command, any UNIX editor or visualizing command (e.g., **vi**, **view**, **more**, **tail**) can be used to review the log file.
 - The log contains entries related to activities and communications associated with actions by AMASS to direct ACSLS robotic activities; the entries should appear in format similar to the following sample:

```
Sep 24 09:49:42 p0drg01 amass LIBSCHED3[7215638]:  
E7003(16)<00000>:xdiStk2749: STK Response received; Status: 0  
  
Sep 24 09:49:42 p0drg01 amass LIBSCHED3[7215638]:  
E7003(16)<00000>:xdiStk2797: ACSLS ACK response received  
  
Sep 24 09:49:42 p0drg01 amass LIBSCHED3[7215638]:  
E7003(16)<00000>:xdiStk2742: Waiting for ACSLS response  
  
Sep 24 09:49:51 p0drg01 amass LIBSCHED3[7215638]:  
E7003(16)<00000>:xdiStk2749: STK Response received; Status: 0
```

```

Sep 24 09:49:51 p0drg01 amass LIBSCHED3[7215638]:
E7003(16)<00000>:xdiStk2873:  ACSLS final response received

Sep 24 09:49:51 p0drg01 amass LIBSCHED3[7215638]:
E7003(16)<00000>:xdiStk2876:  1 network packets transfered

Sep 24 10:18:52 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiArch39:  Archive index : 0

Sep 24 10:18:52 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiStk486:  Media Id = P10011, Drive index = 0

Sep 24 10:18:52 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiStk533:  Sending a mount command

Sep 24 10:19:32 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiArch39:  Archive index : 0

Sep 24 10:19:32 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiStk486:  Media Id = P20676, Drive index = 1

Sep 24 10:19:32 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiStk533:  Sending a mount command

Sep 24 10:34:56 p0drg01 amass LIBSCHED1[7215870]:
E1043(7)<00000>:libsched3165:  Idle Eject timer expired on volume 188
in drive 2.

Sep 24 10:35:07 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiArch39:  Archive index : 0

Sep 24 10:35:07 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiStk686:  Media Id = P20676, Drive index = 1

Sep 24 10:35:07 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiStk719:  Sending a dismount command

Sep 24 10:35:07 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiStk2742:  Waiting for ACSLS response

Sep 24 10:35:07 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiStk2749:  STK Response received; Status: 0

Sep 24 10:35:07 p0drg01 amass LIBSCHED1[7215870]:
E7003(16)<00000>:xdiStk2758:  Error unexpected sequence number: 101 -
expected sequence number: 109

```

- Examine the sections of the log with entries near the time of any problem being investigated, looking for messages that indicate whether there was successful communication between AMASS and ACSLS regarding mounting of a tape and transfer of information. It may be useful to search the log for occurrences of the word **fail** (while viewing the log with **pg**, **view**, **vi**, or other viewing/editing tool, type **/fail** and press the **Return/Enter** key).

- If the log indicates problems in communication between AMASS and ACSLS, it may be useful to use the **quedisplay** command to obtain the AMASS view of the queue and the **medialist** command to obtain the robot view. If these commands show discrepancies indicating a lack of synchrony between AMASS and ACSLS, it may be possible to re-establish that synchrony using the **mediamove** command (refer to procedure **Using mediamove to Establish Synchrony Between quedisplay and medialist** [previous section of this lesson]).
 - **Note:** The message "Error unexpected sequence number: 101 -expected sequence number: 109" is an artifact likely to be removed in releases of AMASS subsequent to Version 5.0.0 Revision 17 and does not reflect a real error.
-

Handling a Data Insertion Failure

Successful data insertion requires interactions among numerous servers, and the interactions are reflected in entries in the debug logs for those servers. Detection and initial isolation of a problem that prevents successful insertion may require tracing events across multiple log files on different hosts. The following procedure is applicable.

Handling a Data Insertion Failure

- 1 At the host for SDSRV (e.g., e0acs05, g0acs03, l0acs03, n0acs04), review the debug log **EcDsScienceDataServerDebug.log** (use procedure **Checking Server Log Files** [previous section of this lesson]).
 - Examine the section of the log with entries near the time of the problem, looking for error messages that indicate communication failure.
 - If the log file entries indicate a communication problem, note the server(s) with which there is impairment or disruption of communication.
- 2 At the host for Archive Server (e.g., e0drg11, g0drg01, l0drg01, n0drg01), review the debug log **EcDsStArchiveServerDebug.log** (use procedure **Checking Server Log Files** [previous section of this lesson]).
 - Examine the section of the log with entries near the time of the problem, looking for error messages that indicate communication failure.
 - If the log file entries indicate a communication problem, note the server(s) with which there is impairment or disruption of communication.

- 3 If Step 1 and/or Step 2 resulted in detection of a problem in the interaction of SDSRV and/or Archive Server with other servers, at the host(s) for those servers, review the server debug log(s). These logs may include:

EcDsStStagingDiskServerDebug.log (on Archive Server host).

EcDsStCacheManagerServerDebug.log (on Archive Server host).

EcDsStRequestManagerServerDebug.log (e.g., on e0dis02, g0dis02, l0dis02, n0dis02; use procedure **Checking the Request Manager Server Debug Log** [previous section of this lesson]).

EcIoAdServerDebug.log (e.g., on e0ins02, g0ins02, l0ins02, n0ins02).

EcSbSubServerDebug.log (e.g., on e0ins01, g0ins01, l0ins01, n0ins01).

- If there is evidence of requests not succeeding or other communication failure, it may be necessary to have System Administrators or Engineering Support personnel resolve the problem (e.g., restart affected servers, execute **EcCsIdPingServers**, ensure that the **Name Server** is up in the mode being used and that its debug log reflects appropriate look-up activity by the application servers, mount points are intact, and database access is not impaired).

Note: The next three steps address running the Check Archive script, **EcDsCheckArchive**. To run this script, it is necessary to enter eight database-specific parameters when prompted during the running of the script: STMGT SQL server name, STMGT database name, STMGT SQL server userID, STMGT SQL server database password, SDSRV SQL server name, SDSRV database name, SDSRV SQL server userID, and SDSRV database password. To facilitate the smooth execution of the script, the parameters may be set as environmental variables instead. The parameters are not readily available to most operators; therefore, you will need to obtain them from the Database Administrator or have the Database Administrator run the script for you, using steps 4 through 6.

- 4 On the host for the Archive Server, type **cd /usr/ecs/<MODE>/CUSTOM /utilities** and then press the **Return/Enter** key.

- The prompt reflects the directory change to **/usr/ecs/<MODE>/CUSTOM/utilities**.

- 5 Type **EcDsCheckArchive <MODE>**.

- The Check Archive script runs; the initially displayed information should be similar to the following:

```
=====
This script is designed to validate the Inventory
against the Archive.

The user must select the menu option associated with the
Volume Group to be validated

Please press [RETURN] to continue
=====
```

- 6 Follow the on-screen prompts for the script, entering the necessary parameters.
 - The script provides indication of any discrepancies between the presence of granules in the Archive and entries in the inventory (metadata). Note that the appearance of a discrepancy is not necessarily indication of a failure (e.g., if a granule has been deleted but the inventory database has not been cleaned up, there may be inventory entries for which there are no granules in the archive), but a problem may be indicated if a discrepancy is apparent for a granule that you just inserted. Note also that this script would not reveal a problem if you attempted to insert a granule which failed to get inserted and also had its metadata fail to be inserted into the inventory (i.e., no granule and no inventory entry = no discrepancy). Therefore, if the script reveals no discrepancies, it may still be useful to conduct a direct examination to determine if the granule has been inserted.
- 7 On the host for the Archive Server, type the directory change command **cd /dss_stk1/<MODE>/<data_type_directory>** and then press the **Return/Enter** key.
 - The working directory is changed to **/dss_stk1/<MODE>/<data_type_directory>**.
- 8 Type **ls -al | grep "<date>"** where "<date>" is a three-letter abbreviation for the month followed by a number indicating the day (e.g., "**Apr 21**") for the granule being inserted, and then press the **Return/Enter** key.
 - If the inserted file is displayed, with date and time of entry, go to Step 9.
 - If the inserted file is not displayed, have the Ingest/Distribution Technician insert the file again. If this succeeds (i.e., the file is now listed), go to Step 9; otherwise, conduct the procedure for **Diagnosing/Investigating Write Errors**, Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, Procedure 17.7.3).
- 9 Determine if the inserted file is reflected in the Inventory Database (Database Administrator function) by logging into Sybase on the host for SDSRV and then selecting the data type for the granule being inserted.
 - If the inserted file is reflected in the Inventory Database, go to Step 10.
 - If the inserted file is not reflected in the Inventory Database, ensure that database access is not impaired (Database Administrator function).
- 10 Determine if the directory from/to which the copy is being made is visible on the machine being used; have the System Administrators or Engineering Support personnel check the mount points on the Archive host and the SDSRV host.
 - If the mount points are OK, go to Step 11.
 - If necessary, have the System Administrators or Engineering Support personnel re-establish the mount point(s).

- 11 If you inserted the file with the DSS Driver, go to Step 13. If you used Ingest to insert the file, on the Ingest host (e.g., e0icg11, g0icg01, l0acg02, n0acg01) examine the **drp**- or **icl**-mounted staging directory to determine if a staging disk was created. To do this, first type **cd /usr/ecs/<MODE>/CUSTOM/drp/<host>/data/staging/cache** (or type **cd /usr/ecs/<MODE>/CUSTOM/icl/<host>/data/StagingArea/cache**), then press the **Return/Enter** key.
- The prompt reflects a change to the specified directory. [*Note:* Be sure that you are checking the correct mount/host. Most ingests use Ingest subsystem staging areas (i.e., **icl**), but others may not. Media ingest (e.g., from tape) typically involves staging in a **dip** area. For a polling ingest for data from EDOS, the polling directory may serve as the staging area. Some data are staged directly to working storage in the Data Server subsystem. If in doubt, consult Ingest/Archive personnel.]
- 12 Type **ls -al | more** and then press the **Return/Enter** key.
- Any staging areas are listed in output similar to the following sample:

```
-rw-rw-r--      1 cmshared cmshared   10375 Jan 30 14:46
:SC:L70RF2.002:16015:6.HDF-EOS
-rw-rw-r--      1 cmshared cmshared   535563 Jan 30 14:46
:SC:L70RF2.002:16015:7.HDF-EOS
-rw-rw-r--      1 cmshared cmshared   154399 Jan 25 12:34
:SC:L7CPF.002:13835:1.ASCII
-rw-rw-r--      1 cmshared cmshared   154399 Jan 25 14:17
:SC:L7CPF.002:16644:1.ASCII
-rw-rw-r--      1 cmshared cmshared   154399 Jan 25 17:31
:SC:L7CPF.002:16769:1.ASCII
-rw-rw-r--      1 cmshared cmshared    67466 Jan 25 18:11
:SC:L7IGS.001:16789:1.ASCII
-rw-rw-r--      1 cmshared cmshared    43570 Jan 25 18:04
:SC:L7IGS.001:16790:1.ASCII
-rw-rw-r--      1 cmshared cmshared 499804704 Feb  6 11:49
:SC:MOD000.001:11856:1.CCSDS
-rw-rw-r--      1 cmshared cmshared 320663592 Feb  6 11:51
:SC:MOD000.001:11856:2.CCSDS
-rw-rw-r--      1 cmshared cmshared     540 Feb  6 11:51
:SC:MOD000.001:11856:3.CCSDS.
```
 - If a staging area for the inserted file appears at the end of the list, go to Step 13.
 - If no staging area appears for the inserted file, it is possible that the ingest failed and that the staging area was immediately removed as part of clean-up. Check the Ingest logs (e.g., **EcInReqMgrDebug.log**, **EcInAutoDebug.log**, **EcInGranDebug.log**, or **EcInGranDebug.log**, depending on the type of Ingest) (refer to procedures for troubleshooting Ingest problems, Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, Chapter 16) to determine if a staging disk was created. If no staging disk was created, it may be necessary to resolve a communications failure as described in Step 7.
- 13 Ensure that the Archive volume groups are set up correctly (refer to procedure **Using Storage Management GUIs to Display Archive Path Information** [previous section of this lesson]).

14 Ensure that the volume groups are on line (refer to procedure **Using *vollist* to Display Volume Data** [previous section of this lesson]).

- If the volume groups are set up correctly and their volumes are on line, and insertion still fails, it is appropriate to contact the Help Desk and prepare a trouble ticket (see Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, Chapter 8).

Handling a Data Acquire Failure

As a first check, it is appropriate to determine if the acquire request appears in the list of System Requests on the Science Data Server GUI. If the acquire request does not appear on the Science Data Server GUI, you will need to determine where the breakdown occurred. Diagnosing an acquire failure requires detailed examination of the following system log files and directories associated with the process:

- Science Data Server log file (EcDsScienceDataServerDebug.log).
- Archive Server log file (EcDsStArchiveServerDebug.log).
- STMGT Request Manager Server log file (EcDsStRequestManagerDebug.log)
- Staging Area.
 - Presence of the relevant file.
 - Staging Disk log files (EcDsStStagingDiskServerDebug.log or EcDsStCacheManagerServerDebug.log).
 - Space available in the staging area.

In addition, note that a number of servers, clients, or other software running on various hosts, as reflected in Table 4, may be involved at various times in processing an acquire request. More information useful in troubleshooting may appear in related logs on these hosts. Use the following procedure for handling a failure to retrieve data.

Table 4. Hosts, Servers, Clients and Other Software Relevant to Acquires

HOST	SERVER/CLIENT/OTHER SOFTWARE
Distribution Server (e.g., e0dis02, g0dis02, l0dis02, n0dis02)	Distribution Server (EcDsDistributionServer) Request Manager Server (EcDsStRequestManagerServer)
Working Storage (e.g., e0wkg01)	Archive Server (EcDsStArchiveServer) Cache Manager Server (EcDsStCacheManagerServer) Staging Disk Server (EcDsStStagingDiskServer)
SDSRV Server (e.g., e0acs05, g0acs03, l0acs03, n0acs04)	Science Data Server (EcDsScienceDataServer) HDF EOS Server (EcDsHdfEosServer)
Access/Process Coordinators (APC) Server (e.g., e0acg11, g0acg01, l0acg02, n0acg01)	Archive Server (EcDsStArchiveServer) FTP Server (EcDsStFtpServer) Cache Manager Server (EcDsStCacheManagerServer) Staging Disk Server (EcDsStStagingDiskServer) Pull Monitor Server (EcDsStPullMonitorServer)
FSMS Server (e.g., e0drg11, g0drg01, l0drg01, n0drg01)	Archive Server (EcDsStArchiveServer) Cache Manager Server (EcDsStCacheManagerServer) Staging Disk Server (EcDsStStagingDiskServer)
Interface Server 02 (e.g., e0ins01, g0ins01, l0ins01, n0ins01)	Subscription Server (EcSbSubServer) Event Server (EcSbEventServer)

Handling a Data Acquire Failure

- 1 Launch the Science Data Server GUI (see procedure 17.3.1 **Launching DSS GUIs Using UNIX Commands** [previous section of this lesson]).
- 2 Click on the **System Requests** tab.
 - The **System Requests** window is displayed.
- 3 Examine the requests displayed in the **System Management Requests** field to determine if SDSRV received the acquire request.
 - If the number of request is large, the **Find** button and field below the **System Management Requests** field may be used to enter and search for information in the request, such as the Requester, or the **Filter . . .** button can be used to launch a **System Management Filter Requests** window to limit the number of entries that appear in the **System Management Requests** field.

- 4 On the SDSRV Server host (e.g., e0acs05, g0acs03, l0acs03, n0acs04), review the server logs **EcDsScienceDataServer.ALOG** and **EcDsScienceDataServerDebug.log** (refer to procedure **Checking Server Log Files** [previous section of this lesson]).

- Examine the section of the log with entries near the time of the problem, looking for messages that indicate whether the relevant file was successfully acquired.
- The **EcDsScienceDataServer.ALOG** file should contain entries identifying the file to be acquired by the ShortName of the corresponding ESDT; entries should be similar to the following:

```
PID : 29168:MsgLink :0 meaningfulname
:DsSrSessionExecuteRequestStart0
Msg: Request ID b5156038-03d3-11d3-8d16-c676e82eaa77:????:
executing:
DsSrRequest (1): DsShSciRequestImp: [ svr: ScienceDS, pri: NORMAL
domain: ]: (DsShSciCommandImp: service: INSERT num parameters: 3
category: Parameters are:
-UnnamedPL[SHORTNAME(AST_L1BT) VERSIONID(001)
--MAINGROUP[SHORTNAME(AST_L1BT) VERSIONID(001)
---
METAFILEGROUP[METADATAFILE(/home/cmops/data/SCAST_L1BT.0011279.met)]
---DATAFILEGROUP[DATAFILE(/home/cmops/data/tahoe-north-middle.MTA)]
---DATAFILEGROUP[DATAFILE(/home/cmops/data/tahoe-north-middle.hdf)]]]
WC)
```

- The **EcDsScienceDataServerDebug.log** file should contain entries regarding the acquire activity. The following types of messages should be included in the log file:

```
05/06/99 12:52:01:
About to execute Statement: exec ProcInsertReqDomain 2205,
"UR:10:DsShESDT
UR:UR:15:DsShSciServerUR:13:[VTC:DSSDSRV]:20:SC:AST_L1BT.001:2201"
05/06/99 12:52:01:
About to execute Statement: ProcInsertAcquireCmd 2206, 2205, 3, null,
null, "tester", "FtpPush", "MAIL", "FILEFORMAT", null, "jrattiga",
"abc123", "tldps04", "/home/jrattiga
/push", null, null
```

- If the ShortName does not appear in the file, with a timestamp corresponding to the time of the attempted acquire, SDSRV may not be running, or may not be communicating with other servers. Have the System Administrator or Operations Controller check to be sure the server is up and, if appropriate, resolve the problem (e.g., restart affected servers, execute **EcCsIdPingServers**, ensure that the **Name Server** is up in the mode being used and that its debug log reflects appropriate look-up activity by the application servers, mount points are intact, and database access is not impaired).
 - If the log file does contain entries for the relevant ShortName, and indicates that two files (the file and its associated metadata file) are being distributed, SDSRV has completed its role in the acquire. Go to the next step.
 - If the ALOG contains the ShortName, and also contains an error showing that the data file time stamp does not match the time stamp required by the acquire, the data file needs to be removed from the Science Data Server and reinserted.
 - This is usually done using a script called DsDbCleanGranules.
- 5 To inspect the Archive Server log and Request Manager Server log for error messages associated with the acquire, on the Archive host (e.g., **e0drg11**, **g0drg01**, **l0drg01**, **n0drg01**), review the respective server logs (**EcDsStArchiveServerDebug.log**, **EcDsStRequestManagerServerDebug.log**); refer to procedure **Checking Server Log Files** (previous section of this lesson) and procedure **A Special Case: Checking the Request Manager Server Debug Log** (previous section of this lesson).
- Examine the sections of the logs with entries near the time of the problem, looking for messages that indicate whether the Request Manager handled the request and whether the Archive Server log shows that the relevant file was successfully acquired.
 - If the logs indicate that the relevant file was successfully acquired, go to the next step.
 - If the file was not successfully acquired, it may be necessary to reboot AMASS (see procedure **Rebooting AMASS** [previous section of this lesson]) and investigate the possibility of read errors (see Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, Procedure 17.7.4 **Diagnosing/Investigating Read Errors**).
- 6 To determine whether the file being acquired (or a link to it) and its associated metadata file arrived in the Data Distribution staging area, on the Distribution Server (e.g., **e0dis02**, **g0dis02**, **l0dis02**, **n0dis02**) type **cd /usr/ecs/<MODE>/CUSTOM/drp/<archivehost>/data/staging/cache** and then press the **Return/Enter** key.
- The working directory is changed to the specified directory.
- 7 Type **ls -lrt** and then press the **Return/Enter** key.

- The contents of the directory are displayed.
- 8 Review the listing to determine whether the relevant file and its metadata file arrived in the staging area.
- The display should contain entries similar to the following:


```
lrwxrwxr-x  1 cmshared cmshared  75 Apr 26 12:52
L7CPF19980518_19980518.01 ->
/usr/ecs/TS1/CUSTOM/drp/raven/data/staging/cache/:SC:L7CPF.001:1427:1
.ASCII
-rw-rw-rw-  1 cmshared cmshared  14802 Apr 26 12:52
SCL7CPF.0011427.met
-rw-rw-r--  1 cmshared cmshared    111 Apr 26 13:01
staging.disk.filename.list
-rw-rw-r--  1 cmshared cmshared  2044 Apr 26 13:01
PACKING.LST.115124935248431
```
 - If the relevant files were not successfully staged, the staging log files may reveal the cause; go to Step 9.
 - If the relevant files were successfully staged, an acquire failure could be a result of problems with related servers or software (see Table 4). Have the System Administrator or Operations Controller ensure that the necessary hosts and servers are up.
- 9 To inspect the Staging Disk log for error messages associated with the acquire, on the APC Server host (e.g., **e0acg11**, **g0acg01**, **l0acg02**, **n0acg01**), review the server logs (e.g., **EcDsStStagingDiskServerDebug.log**; **EcDsStCacheManagerServerDebug.log**); refer to procedure **Checking Server Log Files** (previous section of this lesson).
- Examine the section of each log with entries near the time of the problem, looking for messages that indicate whether the relevant files were successfully staged.
 - If the relevant files were not successfully staged, the cause may be a lack of space in the staging area; go to Step 10.
 - If the relevant files were successfully staged, an acquire failure could be a result of problems with related servers or software (see Table 4). Have the System Administrator or Operations Controller ensure that the necessary hosts and servers are up.
- 10 To check the space available in the staging area, on the Distribution Server (e.g., **e0dis02**, **g0dis02**, **l0dis02**, **n0dis02**) type **cd /usr/ecs/<mode>/CUSTOM/drp/<archivehost>/data** and then press the **Return/Enter** key.
- The working directory is changed to the specified directory.

11 Type **df -k .** (be sure to include the ".") and then press the **Return/Enter** key.

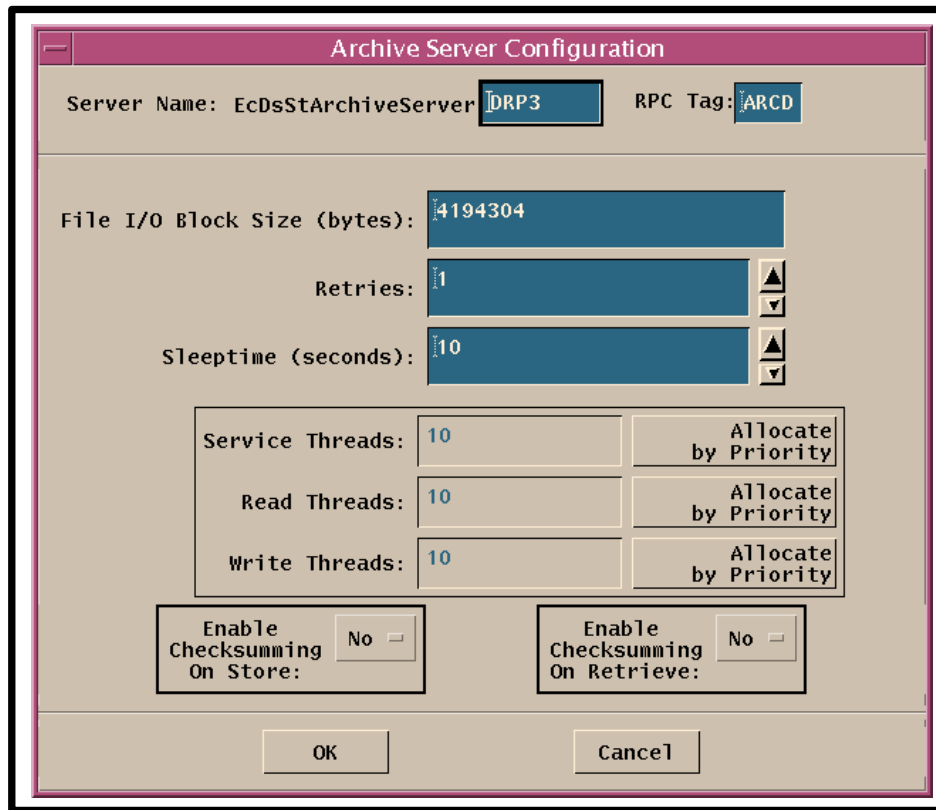
- The filesystem, staging disk space capacity in kbytes, amount used, amount available, and percent of capacity are displayed, as in the following example:

```
Filesystem          kbytes    used    avail capacity  Mounted on
t1drg01:/usr/ecs/TS1/CUSTOM/drp/t1drg01/data
                    225209856 173253056 51956800    77%
/data1/ecs/TS1/CUSTOM/drp/t1drg01/data
```

- If there is not adequate space for staging the relevant files, it will be necessary to free up additional space (e.g., by purging expired files from cache).
-

Checksum De-activation

The system design incorporates calculation of a checksum when a granule is inserted into the archive. If such a checksum is calculated, it can then be used as an indicator to determine if there is data corruption within the archive. Comparison of the original checksum with one calculated, for example, when the granule is retrieved (e.g., for processing or distribution) can detect whether the inserted file and the retrieved file are the same. If the checksums do not match, then the operator can investigate (e.g., by using the **Storage Events** tab of the Storage Management GUI). The checksums are set in the configuration for the archive server, with variables that set calculation on granule insert and calculation on retrieval. The Storage Management GUIs provide an easy way to set these and other STMGT configuration parameters. The settings are available from the **Storage Config** tab, by highlighting the Archive Server and clicking on the **Modify Server** button. This opens the **Archive Server Configuration** window, as illustrated in Figure 29. As the figure shows, the window includes option buttons to **Enable Checksumming On Store:** and **Enable Checksumming On Retrieve:**. Setting checksum calculation variables is addressed in Document 611-CD-610-002 *Mission Operation Procedures for the ECS Project*, Procedure 17.3.5.



The image shows a 'Archive Server Configuration' dialog box. At the top, it has a title bar. Below the title bar, there are two fields: 'Server Name: EcDsStArchiveServer' followed by a blue button labeled 'DRP3', and 'RPC Tag: ARCD' followed by a blue button labeled 'ARCD'. Below these, there are three input fields: 'File I/O Block Size (bytes):' with a value of '4194304', 'Retries:' with a value of '1' and up/down arrow buttons, and 'Sleeptime (seconds):' with a value of '10' and up/down arrow buttons. Below these are three rows of thread configuration: 'Service Threads: 10' with an 'Allocate by Priority' button, 'Read Threads: 10' with an 'Allocate by Priority' button, and 'Write Threads: 10' with an 'Allocate by Priority' button. At the bottom, there are two checkboxes: 'Enable Checksumming On Store:' with a 'No' button, and 'Enable Checksumming On Retrieve:' with a 'No' button. At the very bottom are 'OK' and 'Cancel' buttons.

Figure 29. Archive Server Configuration (from Storage Management GUI)

Calculation of checksums can be time consuming. System throughput may be significantly improved if checksum calculation on granule insert is turned off, and therefore the default reflects checksum calculation turned off. Unfortunately, turning checksums off compromises the ability to detect data corruption in the archive. This problem may be alleviated somewhat by calculating a checksum when a granule is first retrieved from the archive and storing that checksum to be compared with one calculated upon a later retrieval. However, this approach will not guard against the possibility of data corruption on initial insertion (e.g., through I/O read errors).

This page intentionally left blank.

Data Pool Management

The Data Pool is a repository of selected granules with associated metadata and, if available, browse granules, accessible through use of a web browser to search and download by FTP. The Data Pool provides users with a rapid means to obtain granules with associated metadata and any available browse granules. User Services and Science Data Specialists use the Spatial Subscription Server (NSBRV) GUI to create subscriptions for inserting data into the Data Pool and to perform other tasks necessary to manage the Data Pool. Archive and/or engineering support personnel execute functions for Data Pool monitoring, maintenance, and support.

Although users have access to the extensive data holdings of ECS through the normal search and order process, that process may be seen to have some limitations:

- slow distribution times due to tape access.
- search and order tool with a complex user interface.
- limited subsetting options.
- difficult to qualify subscriptions spatially.

The Data Pool provides a means to address these limitations, providing:

- online cache of ECS data.
- easy-to-use drill-down web user interface.
- interface from ECS to external subsetters.
- Spatial Subscription Server.

The data pool is characterized by a large online data cache at each DAAC:

- 67 TB at GSFC (current disk size, usable space 53.6 TB; volume increase likely).
- 49 TB at EDC (current disk size, usable space 39.2 TB; volume increase likely).
- 25 TB at LaRC (current disk size, usable space 20 TB; volume increase likely).
- 2 TB at NSIDC (current disk size, usable space 1.6 TB; volume increase likely).

Data Pool Architecture

Figure 30 provides an overview of the architecture of the Data Pool. The figure illustrates user access to the Data Pool through the web or FTP, and shows use of the Spatial Subscription GUI by an operator, who may be a User Services representative or a Science Data Specialist, to place a Data Pool insertion subscription in the Spatial Subscription Database. As the figure suggests, insertion of a data granule in the ECS archive and associated insertion of metadata in the Science Data Server Database serve to trigger the subscription. The figure also shows Data Pool Maintenance GUI and script tools for operator monitoring and maintenance of the Data Pool.

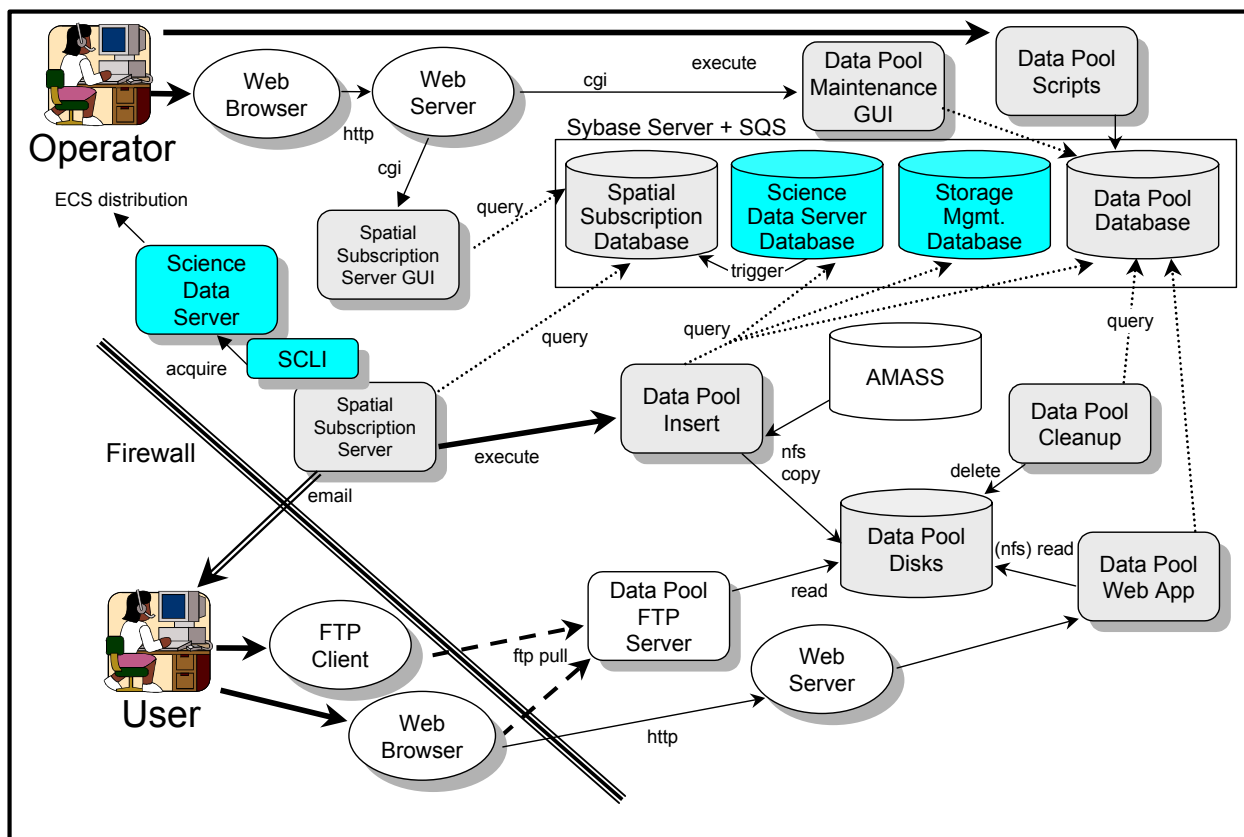


Figure 30. Data Pool Architecture

The limited size of the Data Pool necessitates careful management and maintenance. Data Pool content and data retention are tailored to user needs. Archive and/or engineering support personnel are directly involved in Data Pool monitoring and maintenance, and support User Services and/or Science Data Specialists in managing the content and retention of data in the Data Pool. A major tool used for these functions is the Data Pool Maintenance (DPM) GUI. There are also scripts and utilities for specific maintenance and monitoring functions, and the NSBRV GUI can be used for some maintenance and monitoring functions. Users may contact DAAC operations with a request to place a subscription with ECS to have data inserted in the Data Pool. The subscription is triggered by an ECS event (typically, insertion of a Science Data Granule into the ECS Archive), and as a result a copy of the granule, its associated metadata, and any associated browse granules are inserted into the Data Pool. The subscription also typically results in email notification of the requesting user, and may also result in distribution of data to the user. Subscriptions may be qualified spatially, temporally, or by parameter. Figure 31 provides a more detailed representation of the component architecture and interfaces of the NSBRV.

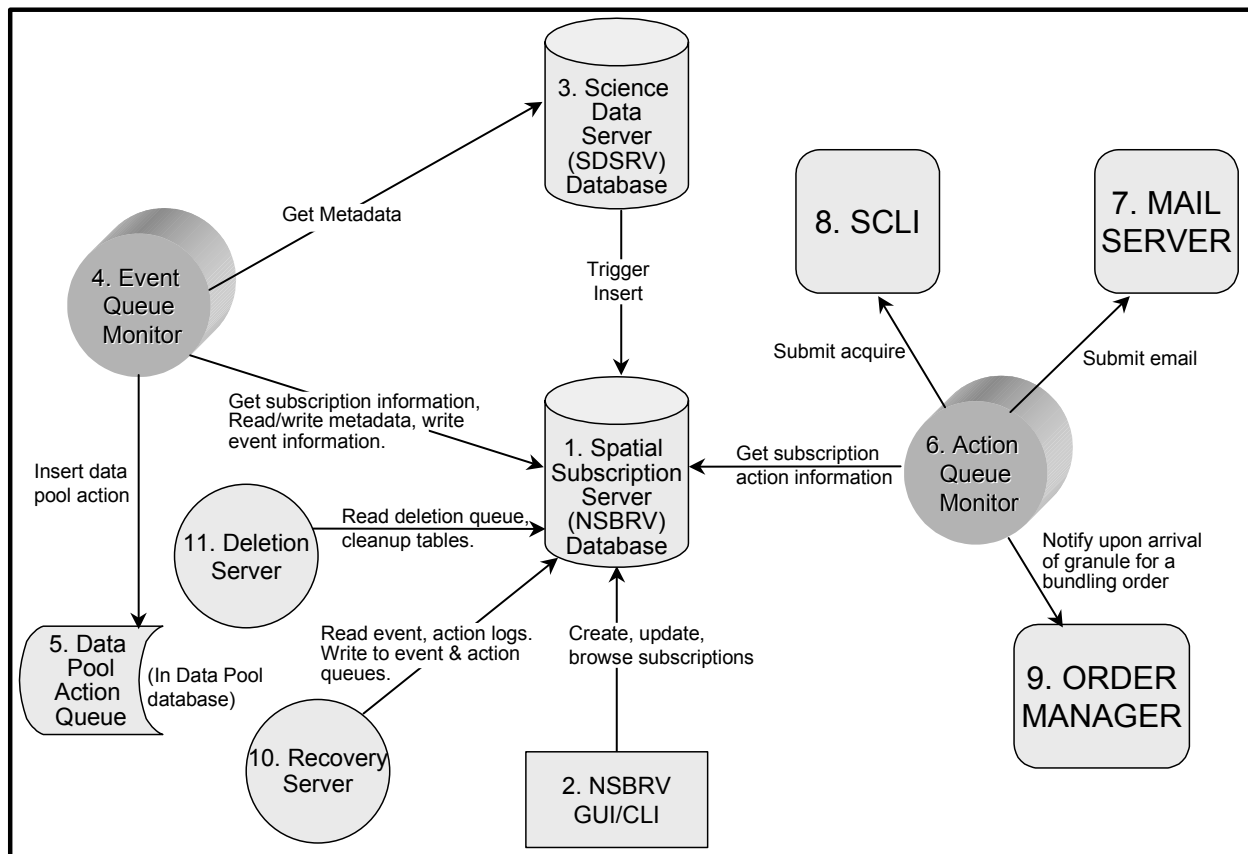


Figure 31. Spatial Subscription Server Component Architecture

The NSBRV database holds subscriptions created with the NSBRV GUI. As the figure suggests, when data are newly inserted into Science Data Server, subscriptions for the data are triggered. This results in placement of an event in the subscribed event queue in the NSBRV database. The Event Queue Monitor (multiple instances of a Perl script) then performs several actions:

- detects and acts on the event.
- gets the metadata for the science data from the Science Data Server.
- stores the metadata in the NSBRV database.
- checks all active subscriptions against the metadata to identify those that qualify for action.
- places qualifying subscriptions into the action queue (a table in the Data Pool database).
- logs its activities in the event queue log.

The Action Queue Monitor (multiple instances of a Perl script) monitors the action queue. It detects events placed there by the Event Queue Monitor and takes appropriate action:

- If the subscription specifies email notification, it composes and sends an email message.
- If the subscription specifies an acquire, it generates acquire parameter files and submits the acquire to SDSRV via the SDSRV Command Line Interface (SCLI).
- If the subscription specifies a bundling order, it sends notification to the Order Manager.

The Action Queue Monitor also logs its activities in the action queue log.

The Recovery Server monitors both the event queue log and the action queue log. It detects events or actions for which processing is stalled. It re-enqueues stalled events in the event queue, and re-enqueues stalled actions in the action queue.

The Deletion Server works off a deletion queue to purge outdated information from the NSBRV database. It purges information related to completed events, including metadata, and it also purges information about completed actions.

Using the Data Pool Maintenance GUI

Most Archive or support personnel tasks for monitoring and maintaining the Data Pool require the use of the Data Pool Maintenance (DPM) GUI. The DPM GUI permits an operator to monitor Data Pool active insert processes, and to check the Data Pool insert queue. It allows an operator to monitor or cancel Data Pool insert actions, and to suspend and resume Data Pool actions. It provides the ability to toggle the state of the NoFreeSpace flag, a flag in the Data Pool database that is set if insertion of a file in the Data Pool fails because of insufficient free space. It may be used to view and update collection groups in the Data Pool database, or to add a collection group to the Data Pool database. The DPM GUI is also used for Data Pool tuning, to configure the number of allowed active insert processes, the default retention period, and the default retention priority.

Other tasks are supported by scripts or utilities. A Data Pool Update Expiration Script (Update Granule Utility) is available for extending the period of retention for selected science granules already in the Data Pool. There is a Data Pool cleanup utility which is typically run in a cron job, but may be invoked manually. Similarly, a utility for accumulating Data Pool access statistics is usually run in a cron job but may be invoked manually. There is a command line utility that permits operators to execute batch inserts of data from the archive into the Data Pool.

Finally, the NSBRV GUI is a major Data Pool management tool. Although used primarily by User Services or science personnel, Archive or engineering support personnel may use it to extend the period of retention in a Data Pool insert subscription, and to view statistics on the processing of events and actions by the NSBRV.

The DPM GUI is a web application; Figure 32 illustrates the DPM GUI **Home Page**, from which the operator can perform some monitoring and maintenance tasks and from which there is access to other pages supporting other tasks.

The DPM GUI Home Page displays the state of several parameters and allows an operator to make changes. It also lists active insert processes. Near the top of the Home page are tabs allowing an operator to access other functions including **Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**. There is also a **Help** tab for assistance in navigation of the GUI.

Data Pool Maintenance

Home Page Batch Summary List Insert Queue Manage Configuration Parameters Manage Collection Groups Manage Themes Help

Refresh Home Page

Screen Refresh Rate (in Secs) Last Screen Refresh on Thu Oct 17 08:43:54 EDT 2002

No Free Space Flag ☒ ON ☐ OFF No Free Space Flag was turned off Oct 3 2002 2:41PM

Suspension ☒ SUSPEND ☐ RESUME Data Pool Insert actions were resumed on Oct 3 2002 2:41PM

Maximum allowed processes 30
Maximum allowed processes from AMASS cache 20
Maximum allowed processes from AMASS tape 3

Total number of active insert processes running 20
Number of active insert processes using AMASS cache 0
Number of active insert processes using AMASS tape 20

Unix ProcessId	ecsID	Collection	Version	Start Time	Status Time	Status	AMASS Cache	Retries
16153	2018651709	MOD28L2	077	Oct 17 2002 6:32AM	Oct 17 2002 6:32AM	COPYING SDSRV INVENTORY	N	0
16180	2018651718	MOD28L2	077	Oct 17 2002 6:32AM	Oct 17 2002 6:33AM	COPYING SDSRV INVENTORY	N	0
16181	2018651722	MOD28L2	077	Oct 17 2002 6:32AM	Oct 17 2002 6:33AM	COPYING SDSRV INVENTORY	N	0
16236	2018651712	MOD28L2	077	Oct 17 2002 6:33AM	Oct 17 2002 6:33AM	COPYING SDSRV INVENTORY	N	0
16309	2018651725	MOD28L2	077	Oct 17 2002 6:33AM	Oct 17 2002 6:33AM	COPYING SDSRV INVENTORY	N	0
16331	2018651759	MOD28L2	077	Oct 17 2002 6:33AM	Oct 17 2002 6:33AM	COPYING SDSRV INVENTORY	N	0

Figure 32. DPM GUI Home Page

Figure 33 illustrates the **Batch Summary** tab. This displays information on inserts made with the command line utility that permits operators to execute batch inserts of data from the archive into the Data Pool. The **Batch Summary** tab displays a summary of status of Data Pool inserts for each batch label. Status includes pending, complete and failed.

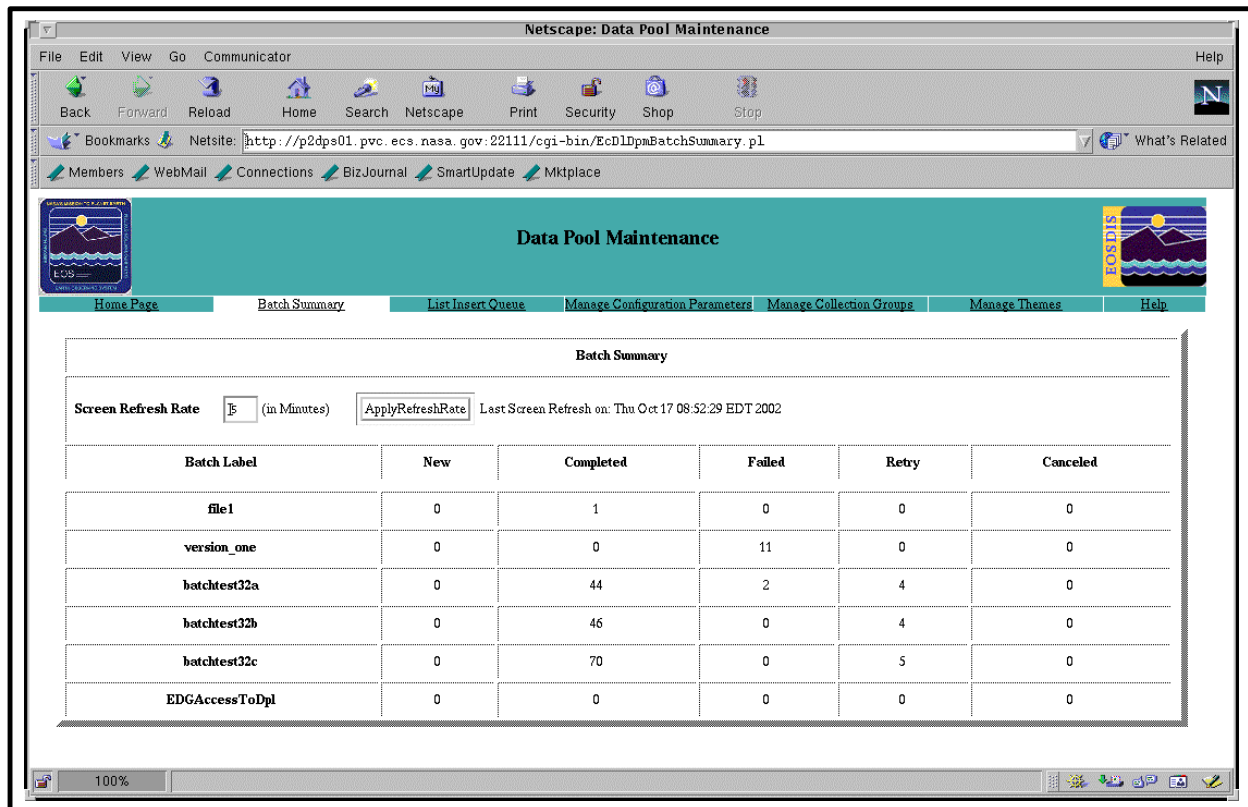


Figure 33. DPM GUI Batch Summary tab

Figure 34 illustrates the **List Insert Queue** tab. This tab provides a list with detailed information on inserts left to process. Insert actions shown in process on the **Home Page** list of active insert processes appear on this list also. The operator can filter the list by choosing a specific batch label from the **Batch Label** pull-down list. The Insert Queue list can also be filtered by Status. For example the operator can choose "Completed" from the **Status** pull-down list and "ALL" from the **Batch Label** pull-down list which will show all the completed inserts for each batch label. After selecting the filter options, click on the **Apply Filter** button to display the filtered list. For each listed insert, a check box in the last column can be used to mark the insert for cancellation. The action to cancel the insert is implemented by a click on an **Apply Change** button at the bottom.

Data Pool Maintenance

Home Page Batch Summary **List Insert Queue** Manage Configuration Parameters Manage Collection Groups Manage Themes Help

Batch Label: Status:

48504

Detailed List of Data Pool Inserts

Click on the NONECS link in the Data Source column to view the XML file associated with this insert

DataSource	BatchLabel	Dispatch Priority	InsertQueueID	SubID	ECSID	Collection	Version	Science Granules and/or Metadata	Enqueue Time	Retries	Status	Click on Box to Cancel
ECS	--	--	20000000039179	529	2018653209	MOD28QC	077	science and metadata	Oct 16 2002 10:36PM	1	RETRY	<input type="checkbox"/>
ECS	--	--	20000000039112	595	2018653154	MYD01	077	science and metadata	Oct 16 2002 10:33PM	1	RETRY	<input type="checkbox"/>
ECS	--	--	20000000039091	597	2018653043	MYD021KM	077	science and metadata	Oct 16 2002 10:30PM	1	RETRY	<input type="checkbox"/>
NONECS	batchtest32b	200	20000000027685	-1	--	AVHRR	002	science and metadata	Oct 16 2002 2:07PM	4	RETRY	<input type="checkbox"/>
NONECS	batchtest32b	200	20000000027690	-1	--	AVHRR	002	science and metadata	Oct 16 2002 2:07PM	4	RETRY	<input type="checkbox"/>
NONECS	batchtest32b	200	20000000027705	-1	--	CZCS	002	science and metadata	Oct 16 2002 2:07PM	4	RETRY	<input type="checkbox"/>
NONECS	batchtest32b	200	20000000027718	-1	--	CZCS	002	science and metadata	Oct 16 2002 2:07PM	4	RETRY	<input type="checkbox"/>

[Continue](#)

Figure 34. DPM GUI List Insert Queue tab

[illegible]

625-CD-610-003

Figure 36 shows the **Manage Configuration Parameters** tab of the DPM GUI. This tab lists numerous Data Pool configuration parameters with their settings and a brief description of each. For each parameter, there is also a check box permitting an operator to modify the parameter. At the bottom of the screen is an **Apply Change** button to implement the change(s).

Data Pool Maintenance

Home Page Batch Summary List Insert Queue **Manage Configuration Parameters** Manage Collection Groups Manage Tables Help

Parameter Name	Parameter Value	Click on Box to Modify Param
ActionQueueCleanupFrequency	300 Frequency in seconds when the action queue is checked for completed actions and those older than the configured retention are removed.	<input type="checkbox"/>
BatchSummaryAutoRefresh	5 Auto Refresh Rate for Batch Summary Screen	<input type="checkbox"/>
DefaultRetentionPeriod	30 The number of days is configured for the default retention period with all Data Pool Insert actions.	<input type="checkbox"/>
DefaultRetentionPriority	300 The default retention priority is configured for the Data Pool Insert action. The value should be between 1 and 255.	<input type="checkbox"/>
DeleteCompletedActionsAfter	3600 The time in minutes that operators let completed actions stay in the insert action queue before making them eligible for removal.	<input type="checkbox"/>
HEGCleanupAge	30 HEG Cleanup Age in days	<input type="checkbox"/>
IdleSleep	5 The number of seconds to sleep when there is nothing to do.	<input type="checkbox"/>
InCacheTimeLimit	30 Max time in minutes operators are willing to wait for a DPU process to complete whose files are in cache	<input type="checkbox"/>
InsertRetryWait	30 The number of seconds to wait before a insert that failed reliably should be resubmitted.	<input type="checkbox"/>
MaxInsertRetries	3 The maximum number of times an insert should be tried again (-1 means forever).	<input type="checkbox"/>
NewActionCheckFrequency	30 Frequency in seconds for checking new actions.	<input type="checkbox"/>
NumOAllowsCacheProcesses	20 Max number of insert processes that require AMASS access to cache	<input type="checkbox"/>
NumOAllowsInsertProcesses	30 Max number of insert processes	<input type="checkbox"/>
NumOAllowsInCacheProcesses	10 Max number of insert processes that require AMASS access to tape	<input type="checkbox"/>
OnTapeTimeLimit	30 Max time in hours that operators are willing to wait for DPU process to complete whose files are not in cache	<input type="checkbox"/>
RefreshRate	30 Refresh rate in seconds when the list of active insert process Data Maintenance Page is refresh.	<input type="checkbox"/>
RunAwayCheckFrequency	30 Frequency in seconds for checking run away processes.	<input type="checkbox"/>
RunawayDuration	3600 Max period of time to wait for an insert to complete.	<input type="checkbox"/>
SizeOfInsertQueueList	500 Insert Queue List chunk size	<input type="checkbox"/>
StartUpWait	30 The number of seconds to delay start-up while trying to clean out left over DPU processes.	<input type="checkbox"/>

Apply Change

Figure 36. DPM GUI Manage Configuration Parameters tab

Figure 37 illustrates the **Manage Collection Groups** tab. This tab lists the collection groups, providing for each the Data Source (ECS or NONECS), a Group ID, and a brief description of the collection. Each Group ID is a link. At the bottom of the page, there are links permitting an operator to **Add a Collection Group**, or **Modify a Collection Group Description**. A click on one of the Group ID links brings up a page listing the collection(s) in that group along with links to add or modify a collection.

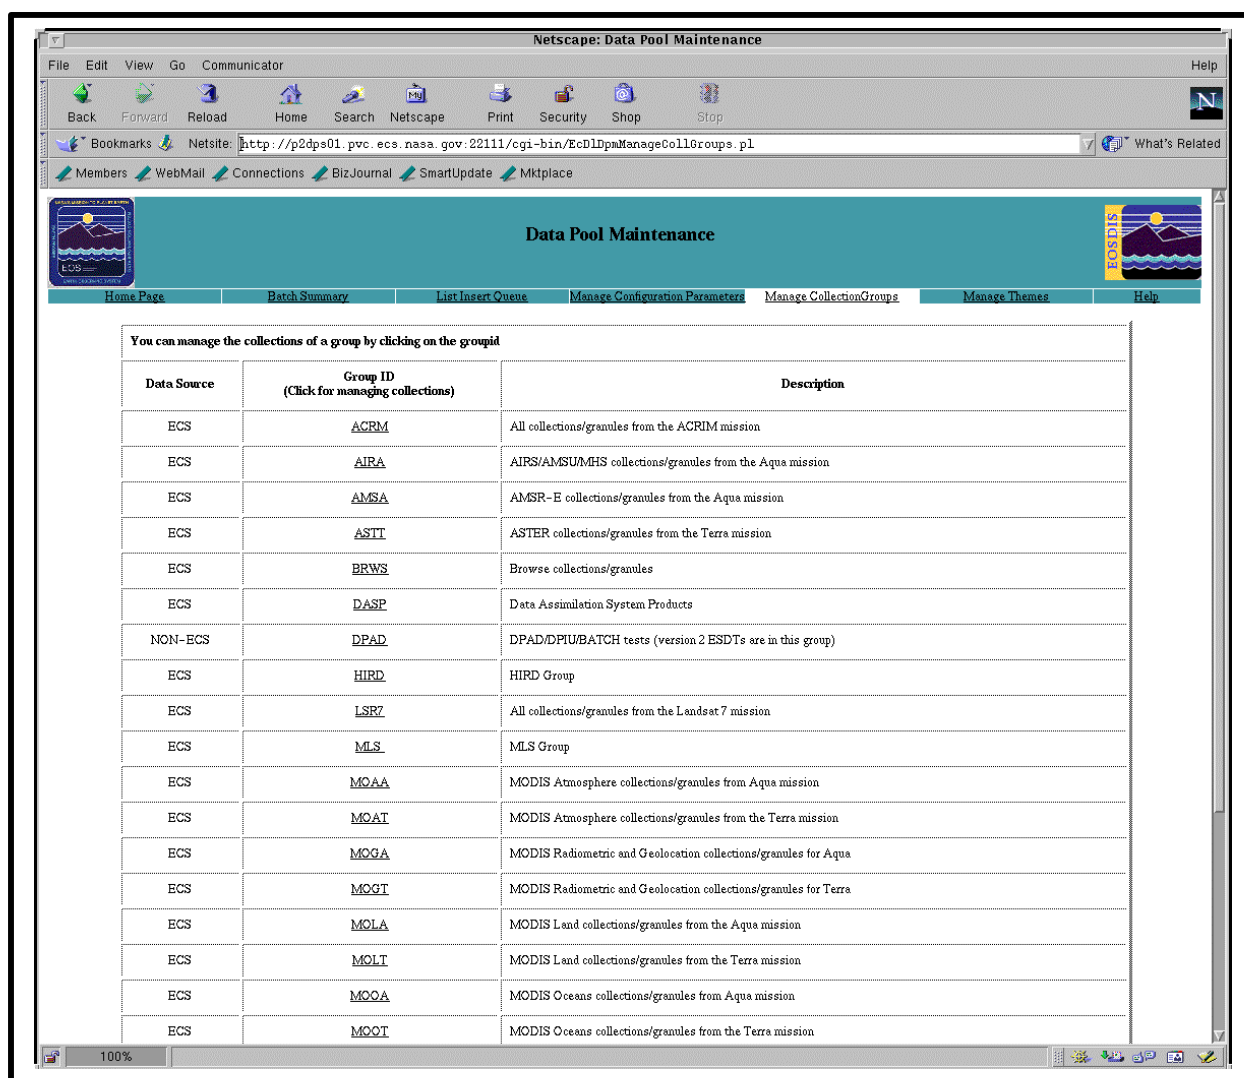


Figure 37. DPM GUI Manage Collection Groups tab

Figure 38 shows the appearance of the **List of Collections** screen obtained by clicking on one of the Group ID links on the **Manage Collection Groups** page. Clicking on a link for a NONECS collection or for an ECS collection returns the appropriate NONECS or ECS list, respectively. On the resulting screen shown, the entries in the **Collection** column are links, and a click on one of them results in display of a **Description for Collection** screen, also shown in the figure. These screens provide a means of determining what collections within a collection group have been designated valid for Data Pool insertion and whether the insertion is for science granules and metadata or metadata only.

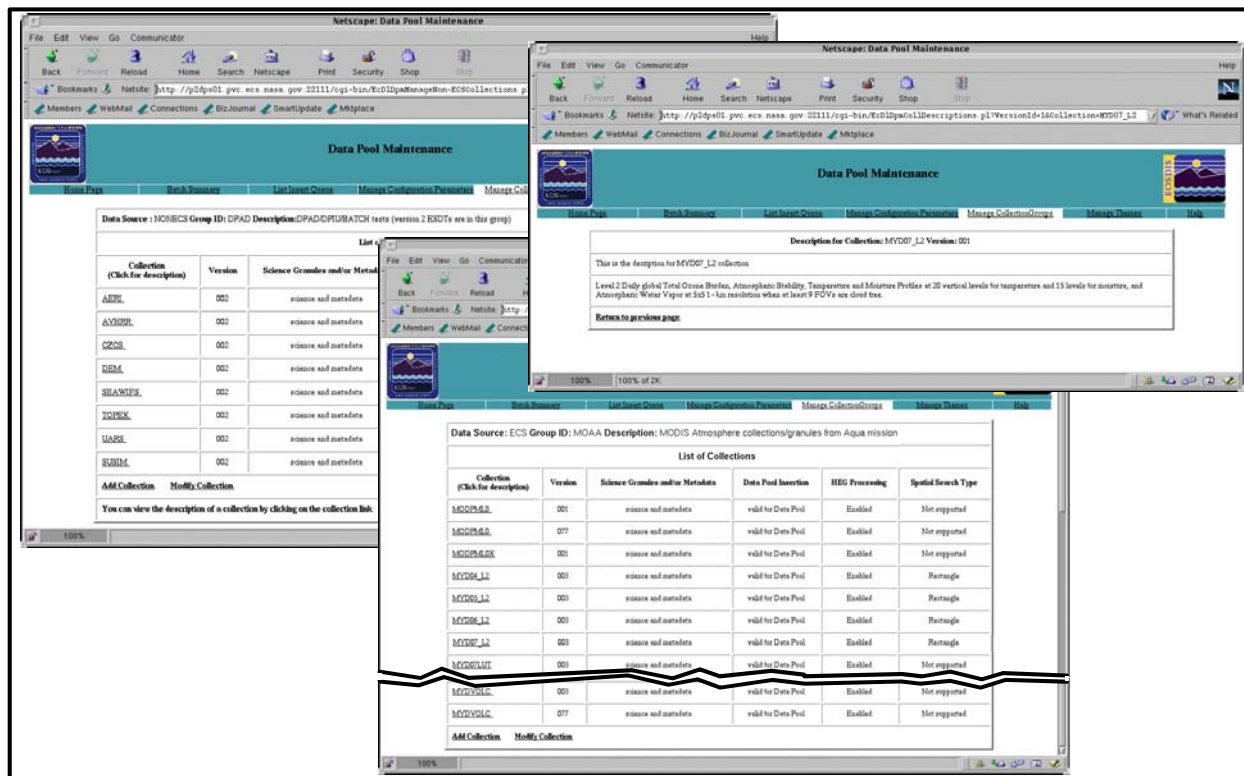


Figure 38. DPM GUI List of Collections and Description for Collection screens

Figure 39 shows the screen obtained by clicking on the **Modify Collection Group Description** link at the bottom of the **Manage Collection Groups** screen. On this screen, the operator can modify the description of one or more collection groups, then mark the group for change, and implement the change with a click on the **Apply Change** button at the bottom.

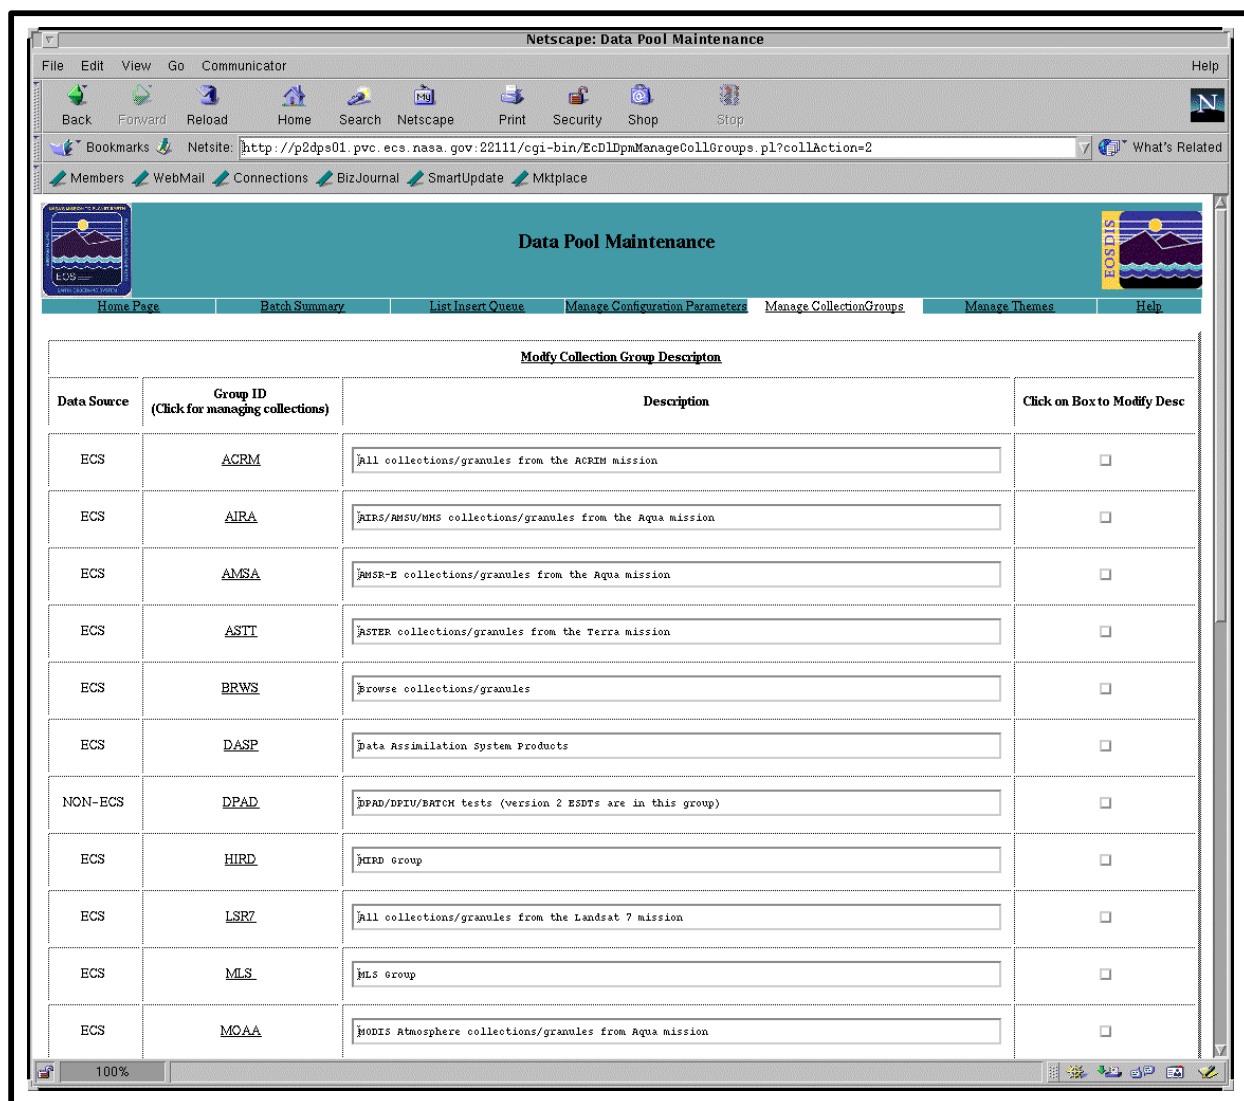


Figure 39. DPM GUI Modify Collection Group Description screen

Figure 40 shows the screen obtained by clicking on the **Add Collection Group** link at the bottom of the **Manage Collection Groups** screen. On this screen, the operator can use an option button to specify whether the new collection group is for ECS or NONECS collections, enter a Group ID, and provide a description for one or more collection groups. The then adds the group(s) with a click on the **Apply Change** button at the bottom.

Caution:

The Add Collection Group function is to be exercised judiciously because the DPM GUI does not provide any means of deleting collection groups.

Figure 40. DPM GUI Add Collection screen

Recall that the **List of Collections** screens (illustrated in Figure 38) have links at the bottom providing access to screens for adding or modifying collections. Figure 41 illustrates the screen obtained by clicking on the **Modify Collection** link on the screen from Figure 38 for ECS Collections. This screen is organized in a format similar to the **List of Collections** screen, except that the entries in the **Science Granules and/or Metadata** column and **Data Pool Insertion** column are option buttons permitting the operator to select **science and metadata** or **metadata only** in one column and **valid for data pool** or **invalid for data pool** in the other. There is also a column with check boxes to mark the collection(s) for change, and an **Apply Change** button to implement the marked change(s).

Data Source : ECS Group ID: MOAT **Description:** MODIS Atmosphere collections/granules from the Terra mission

Modify Collections

You can modify the science or metadata exclusion flag and/or data pool insertion flag by choosing your selection from the drop down lists. After selecting check the corresponding check box for action

Collection	Version	Science Granules and/or Metadata	Data Pool Insertion	Click on Box to Modify Collection
MOD000	001	science and metadata <input type="button" value=""/>	valid for data pool <input type="button" value=""/>	<input type="checkbox"/>
MOD000	077	science and metadata <input type="button" value=""/>	valid for data pool <input type="button" value=""/>	<input type="checkbox"/>
MOD000X	001	science and metadata <input type="button" value=""/>	valid for data pool <input type="button" value=""/>	<input type="checkbox"/>
MOD000X	077	science and metadata <input type="button" value=""/>	valid for data pool <input type="button" value=""/>	<input type="checkbox"/>
MOD001	001	science and metadata <input type="button" value=""/>	valid for data pool <input type="button" value=""/>	<input type="checkbox"/>
MOD041_O	001	science and metadata <input type="button" value=""/>	valid for data pool <input type="button" value=""/>	<input type="checkbox"/>
MOD04_L2	001	science and metadata <input type="button" value=""/>	valid for data pool <input type="button" value=""/>	<input type="checkbox"/>
MOD04_L2	003	science and metadata <input type="button" value=""/>	valid for data pool <input type="button" value=""/>	<input type="checkbox"/>
MOD04_L2	077	science and metadata <input type="button" value=""/>	valid for data pool <input type="button" value=""/>	<input type="checkbox"/>

Figure 41. DPM GUI Modify ECS Collection screen

Figure 42 illustrates the screen obtained by clicking on the **Modify Collection** link on the screen from Figure 38 for NONECS Collections. This screen permits the operator to make changes like those available for modifying ECS Collections. For NONECS collections, the operator also has the option to modify the Spatial Search Type (options include "Orbit," "Gpolygon," "Rectangle," and "Not supported"). Spatial Search Type cannot be modified if granules have been inserted for the collection or the collection is enabled for insert. It can only be changed if the collection is invalid for insert and there are no granules associated with the collection.

Data Source : NONECS Group ID: DPAD Description: DPAD/DPIUBATCH tests (version 2 ESDTs are in this group)

Modify Collections

You can modify the Science Granule exclusion flag and (or) Data Pool Insertion flag and (or) Spatial Search Type flag by choosing your selection from the drop down lists. After selecting check the corresponding check box for action

Collection (Click for description)	Version	Science Granules and/or Metadata	Data Pool Insertion	Spatial Search Type	Click on Box to Modify Collection
AERI	002	science and metadata <input type="checkbox"/>	valid for data pool <input type="checkbox"/>	Rectangle <input type="checkbox"/>	<input type="checkbox"/>
AVHRR	002	science and metadata <input type="checkbox"/>	valid for data pool <input type="checkbox"/>	Gpolygon <input type="checkbox"/>	<input type="checkbox"/>
CZCS	002	science and metadata <input type="checkbox"/>	valid for data pool <input type="checkbox"/>	Gpolygon <input type="checkbox"/>	<input type="checkbox"/>
DEM	002	science and metadata <input type="checkbox"/>	valid for data pool <input type="checkbox"/>	Rectangle <input type="checkbox"/>	<input type="checkbox"/>
SEAWIFS	002	science and metadata <input type="checkbox"/>	valid for data pool <input type="checkbox"/>	Rectangle <input type="checkbox"/>	<input type="checkbox"/>
TOPEX	002	science and metadata <input type="checkbox"/>	valid for data pool <input type="checkbox"/>	Gpolygon <input type="checkbox"/>	<input type="checkbox"/>
UARS	002	science and metadata <input type="checkbox"/>	valid for data pool <input type="checkbox"/>	Not supported <input type="checkbox"/>	<input type="checkbox"/>
SUSIM	002	science and metadata <input type="checkbox"/>	valid for data pool <input type="checkbox"/>	Not supported <input type="checkbox"/>	<input type="checkbox"/>

[Return to previous page](#)

Figure 42. DPM GUI Modify NONECS Collection screen

The **Add Collection** screen for ECS Collections, illustrated in Figure 43, is organized in the same way and works the same way as the **Modify Collection** screen. The collections listed in the **Collection** column, however, are obtained from the Science Data Server database, representing those collections in the database not currently set as valid for insertion into the Data Pool.

Collection	Version	Science Granules and/or Metadata	Data Pool Insertion	Click on Box to Add Collection
TAST_04	001	science and metadata <input type="checkbox"/>	invalid for data pool <input type="checkbox"/>	<input type="checkbox"/>
TAST_08	001	science and metadata <input type="checkbox"/>	invalid for data pool <input type="checkbox"/>	<input type="checkbox"/>
TAST_L1B	001	science and metadata <input type="checkbox"/>	invalid for data pool <input type="checkbox"/>	<input type="checkbox"/>
TAST_05	001	science and metadata <input type="checkbox"/>	invalid for data pool <input type="checkbox"/>	<input type="checkbox"/>
MOD02QC	001	science and metadata <input type="checkbox"/>	invalid for data pool <input type="checkbox"/>	<input type="checkbox"/>
MOD03CP	001	science and metadata <input type="checkbox"/>	invalid for data pool <input type="checkbox"/>	<input type="checkbox"/>
MOD03QC	001	science and metadata <input type="checkbox"/>	invalid for data pool <input type="checkbox"/>	<input type="checkbox"/>
TEST04	001	science and metadata <input type="checkbox"/>	invalid for data pool <input type="checkbox"/>	<input type="checkbox"/>
AIRBOSCI	001	science and metadata <input type="checkbox"/>	invalid for data pool <input type="checkbox"/>	<input type="checkbox"/>

Figure 43. DPM GUI Add Collection screen, ECS Collections

The **Add Collection** screen for NONECS Collections, illustrated in Figure 44, is organized in similar fashion, but the operator must enter **Collection** and **Version** information in fields provided for the purpose. It is also necessary to specify **Spatial Search Type**. There are four options for Spatial Search Type: "Orbit", "Gpolygon", "Rectangle", and "Not supported." The operator can set the value by choosing one of them. "Not Supported" is the default.

Data Pool Maintenance

Home Page Batch Summary List Insert Queue Manage Configuration Parameters Manage Collection Groups Manage Themes Help

Data Source: NONECS Group ID: DPAD Description: DPAD/DPIU/BATCH tests (version 2 ESDTs are in this group)

Add Collections

Collection	Version	Science Granules and/or Metadata	Data Pool Insertion	Spatial Search Type
<input type="text"/>	<input type="text"/>	Science and Metadata <input type="checkbox"/>	Valid for data pool <input type="checkbox"/>	Not supported <input type="checkbox"/>

[Return to previous page](#)

Figure 44. DPM GUI Add Collection screen, NONECS Collections

Figure 45 illustrates the **Manage Themes** tab. This tab allows the operator to view a list of themes in alphabetical order. The list can be filtered using the pull-down list and choosing one of four options: Beginning Letters, Insert Enabled, Web Enabled, and NONE. After selecting the options, a click on the **Apply Filter** button displays the filtered list of themes. The operator can also delete a theme by selecting the corresponding check box and clicking on the **Apply Change** button. There are **Add New Theme** and **Modify Theme** links providing access to screens for managing those functions. After the operator completes adding a new theme or modifying a theme by clicking on the **Apply Change** button at the screens for those functions, the changes take effect in the Data Pool database and the changes are also reflected in the **Manage Themes** page shown in Figure 45.

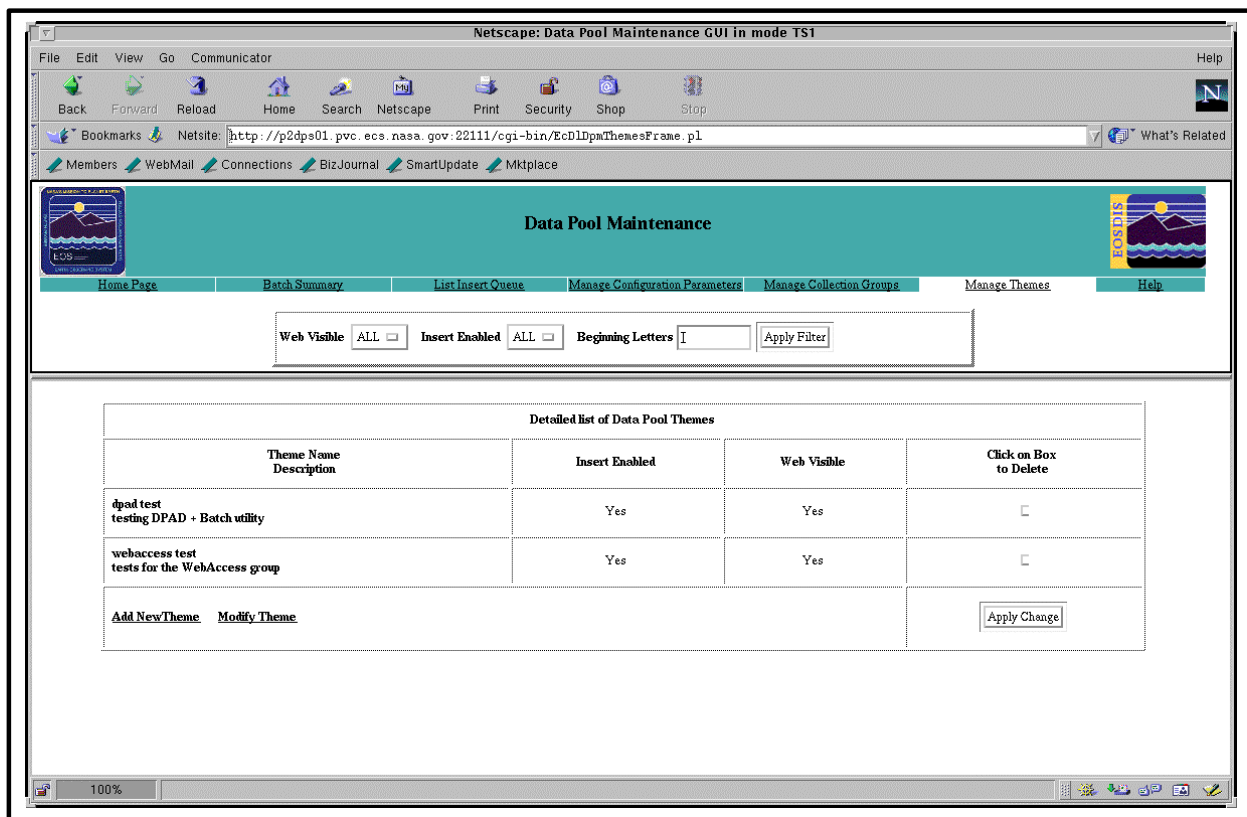


Figure 45. DPM GUI Manage Themes tab

A click on the **Add New Theme** link of the screen shown in Figure 45 results in display of the page shown in Figure 46. To specify a theme, the operator adds information in fields provided for the purpose. **Theme Name** and **Description** are text entry fields. **Insert Enabled** and **Web Enabled** are check boxes to specify whether the theme is valid for Data Pool insert or not and whether it is valid for web drill-down access or not, respectively. A click on the **Apply Change** button commits the changes to the Data Pool database and updates the **Manage Themes** page shown in Figure 45.

Data Pool Maintenance

Home Page | Batch Summary | List Insert Queue | Manage Configuration Parameters | Manage Collection Groups | Manage Themes | Help

Web Visible: ☐ ALL ☐ Insert Enabled: ☐ ALL ☐ Beginning Letters:

Add New Theme

Theme Name	Description	Web Visible	Insert Enabled
Theme: <input type="text"/>	Description: <input type="text"/>	<input type="checkbox"/>	<input type="checkbox"/>

[Return to theme list](#)

Figure 46. DPM GUI Add New Theme page

A click on the **Modify Theme** link of the screen shown in Figure 45 results in display of the page shown in . **Theme Name** is the only field that is not editable. The operator can modify the description of a theme by simply retyping in the text area. The operator also can change the option for **Insert Enabled** and **Web Enabled** by selecting or deselecting the appropriate boxes. A click on the **Apply Change** button commits the changes to the Data Pool database and updates the **Manage Themes** page shown in Figure 45.

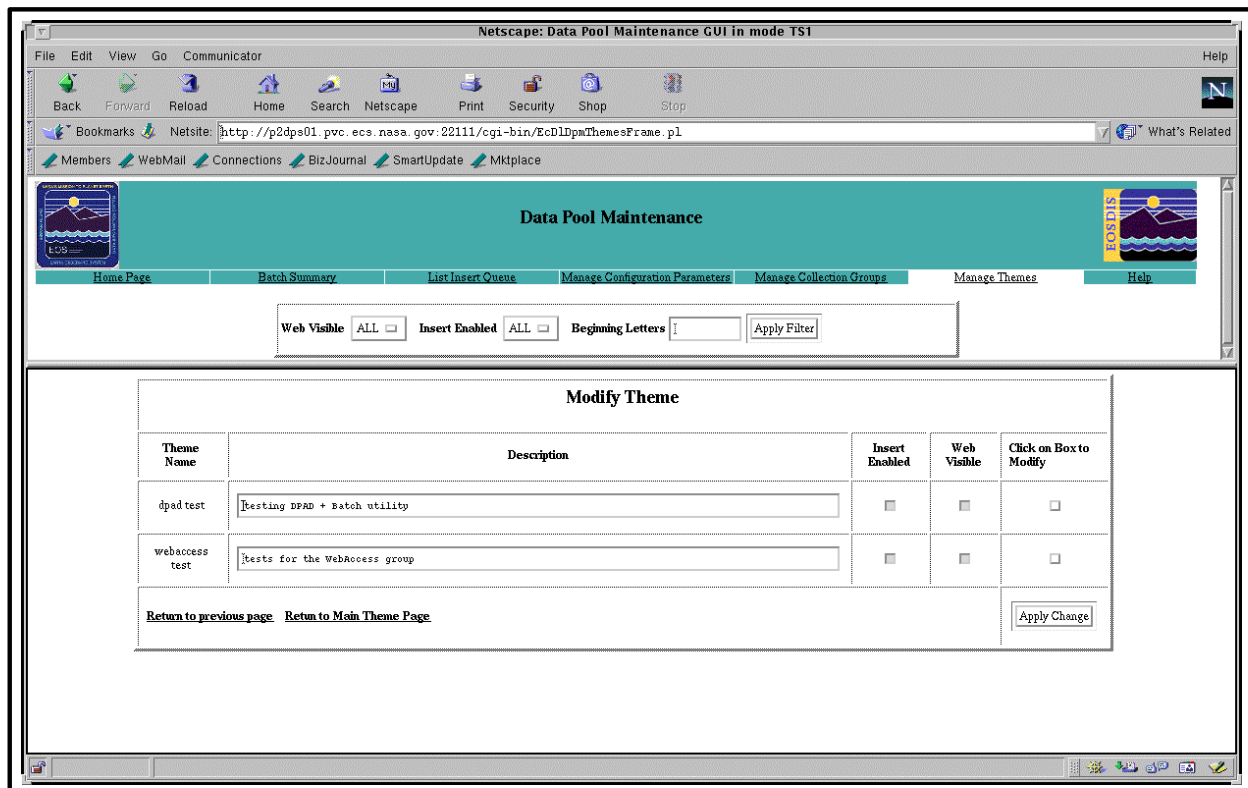


Figure 47. DPM GUI Modify Theme page

Let's examine how you can use the DPM GUI for Data Pool maintenance tasks. The first thing you will need to do for many functions is to launch the GUI. The procedure for launching the GUI is provided separately here and referenced in other procedures.

Launch the DPM GUI

- 1 At the UNIX command shell prompt, type **setenv DISPLAY clientname:0.0** and then press the **Return/Enter** key.
 - For *clientname*, use either the local terminal/workstation IP address or its machine name.

- 2 Start the log-in to a Netscape host by typing **/tools/bin/ssh *hostname*** (*e.g.*, g0ins02, e0ins02, l0ins02, n0ins02) at the UNIX command shell prompt, and press the **Return/Enter** key.
 - If you receive the message, **Host key not found from the list of known hosts. Are you sure you want to continue connecting (yes/no)?** type **yes** (“y” alone does not work).
 - If you have previously set up a secure shell passphrase and executed **sshremote**, a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears; continue with Step 3.
 - If you have not previously set up a secure shell passphrase; go to Step 4.
 - 3 If a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears, type your **Passphrase** and then press the **Return/Enter** key. Go to Step 5.
 - 4 At the **<user@remotehost>'s password:** prompt, type your **Password** and then press the **Return/Enter** key.
 - You are logged in and a UNIX command shell prompt is displayed.
 - 5 Type **netscape** and then press the **Return/Enter** key.
 - The Netscape web browser is displayed.
 - 6 Click in the **Netsite:** field.
 - The field is highlighted.
 - 7 Type the Universal Resource Locator (URL) for the DPM GUI and then press the **Return/Enter** key.
 - The DPM **Home Page** is displayed, offering access to Data Pool maintenance functions (**Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**) as well as a tab for **Help** in navigating the GUI.
-

Monitoring Data Pool Active Insert Processes and Insert Actions

You may wish to keep an instance of the DPM GUI displayed to monitor Data Pool Active Insert Processes. The following procedure is applicable.

Use the DPM GUI to Monitor Data Pool Active Insert Processes

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, providing a table of active insert processes showing columns of detailed information for each process, including:
 - the UNIX process identifier (**Unix ProcessId**).
 - the ECS identifier (**ECS_ID**), or Granule ID for the granule being processed.
 - the **Collection** to which the granule belongs.
 - the **Version** for the collection to which the granule belongs.
 - the time at which the insert processing started (**StartTime**).
 - the time at which the status listed in the next column was achieved (**StatusTime**).
 - the current state of the insert process (**Status**).
 - the **AMASS Cache** availability (**Y** or **N**) of the granule being processed (*Note*: The system is designed for rapid insertion of data into the Data Pool by quickly processing data that are available in cache, such as data that are staged for archiving. If the insert processing is delayed and the data are removed from cache, the Data Pool insert is likely to fail.)
 - the number of attempts (**Retries**) by the process to recover from retrievable errors (e.g., Data Pool disk temporarily unavailable, Data Pool directory does not exist, Data Pool database temporarily unavailable).
- 2 If it is desired to obtain an immediate screen refresh, click on the **Refresh** button near the upper right corner of the display to refresh the screen.
 - The displayed insert action data are updated. (*Note*: The screen refreshes automatically at intervals determined by the number of seconds specified in the **Screen Refresh Rate** field.
- 3 If it is desired to change the automatic screen refresh rate, click at the end of the **Screen Refresh Rate** field.
 - The cursor is displayed at the end of the **Screen Refresh Rate** field.

- 4 To set a new value for **Screen Refresh Rate**, use the **Backspace** key to remove the current value, type the new value (in seconds), and then click on the **Apply** button to the right of the **Screen Refresh Rate** field.
 - The typed value appears in the **Screen Refresh Rate** field and the automatic refresh frequency is based on the new value.
-

Checking the Status of Batch Inserts

The DPM GUI provides a page to display a summary of the status of batch Data Pool inserts made using the Synergy III batch insert utility. The following procedure is applicable.

Check the Status of Batch Inserts

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, providing status and capabilities for changing the **Screen Refresh Rate**, **NoFreeSpaceFlag**, and **Suspension** state, as well as a table of active insert processes.
 - 2 To view the status of batch insert processes, click on the **Batch Summary** tab.
 - The GUI displays the **Batch Summary** page, providing for each batch label the numbers of inserts for that label that are **New**, **Completed**, **Failed**, in **Retry**, and **Canceled**.
 - The page also shows the screen refresh rate in minutes; the rate may be changed by clicking in the **Screen Refresh Rate** field, replacing the displayed value with the desired value, and clicking on the **Apply Refresh Rate** button.
-

Suspending and Resuming Data Pool Insert Actions

At times, it may be necessary to suspend Data Pool insert actions (e.g., for cleanup or other maintenance activities). The following procedure is applicable.

Suspend and Resume Data Pool Insert Actions

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, providing status and capabilities for changing the **Screen Refresh Rate**, **NoFreeSpaceFlag**, and **Suspension** state, as well as a table of active insert processes.

- 2 To suspend Data Pool insert actions, click the **SUSPEND** selection/indicator button.
 - The selected button is filled to indicate its selection.
 - 3 Click on the **Apply** button to the right of the **SUSPEND** and **RESUME** selection/indicator buttons.
 - The screen is refreshed and the **Suspension** status line shows "Data Pool Insert actions were suspended on <date time>" to indicate the change.
 - 4 To resume Data Pool insert actions, click the **RESUME** selection/indicator button.
 - The selected button is filled to indicate its selection.
 - 5 Click on the **Apply** button to the right of the **SUSPEND** and **RESUME** selection/indicator buttons.
 - The screen is refreshed and the **Suspension** status line shows "Data Pool Insert actions were resumed on <date time>" to indicate the change.
-

Checking the Data Pool Insert Queue and Canceling a Data Pool Insert Action

The **List Insert Queue** tab of the DPM GUI provides a list of Data Pool inserts left to process. It also provides for each listed insert a check box permitting the insert to be marked for cancellation, and an **Apply Change** button to implement the cancellation. If an insert is to be canceled, the following procedure is applicable.

Check the Data Pool Insert Queue and Cancel a Data Pool Insert Action

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**).
- 2 Click on the **List Insert Queue** tab.
 - The **List Insert Queue** page is displayed, providing a table of inserts left to process showing columns of detailed information for each process, including:
 - the insert queue identifier (**InsertQueueID**).
 - the subscription identifier (**SubID**) of the subscription selected by the software for processing (**Note**: There may be multiple subscriptions specifying insertion of specific data into the Data Pool, but only one insert is needed and therefore only one of the subscriptions serves as the basis for the insert action. The **SubID** is of no particular significance to an operator and may safely be ignored.)

- the database identifier (**dbID**), or Granule ID for the granule to be processed.
 - the **Collection** to which the granule belongs.
 - the **Version** for the collection to which the granule belongs.
 - an indication of whether the insert is to include science granules and metadata or just the metadata (**Science Granules and/or Metadata**).
 - the time at which the insert was placed in the insert queue (**Enqueue Time**).
 - the number of attempts (**Retries**) by the process to recover from retrievable errors (e.g., Data Pool disk temporarily unavailable, Data Pool directory does not exist, Data Pool database temporarily unavailable).
- There is a **Continue** link at the bottom of the page; if there are more inserts than can be displayed in the space of one page, this link permits display of the continuation of the list.
- 3** If it is desired to cancel one or more inserts, click on the check box at the end of the row of information for the insert(s) to be canceled.
- The check box for each selected insert is filled to indicate selection.
- 4** To implement the cancellation of any selected insert(s), click on the **Apply Change** button.
- A confirmation message asks "Are you ready to cancel the insert for . . ." and there are links displayed for **Yes, cancel insert** and **No, return to previous page**.
- 5** To confirm cancellation, click on the **Yes, cancel insert** link.
- The **List Insert Queue** page is displayed with the canceled insert(s) removed and the count of inserts left to process reduced by the number of inserts canceled.
-

Working with the NoFreeSpaceFlag

The **NoFreeSpaceFlag** must be **OFF** in order to make inserts into the Data Pool. When an insert requires more space than is available on the Data Pool disks, the insert fails and the **NoFreeSpaceFlag** is set to **ON**. To enable further inserts, it is necessary to free up additional space (e.g., run the Data Pool Cleanup utility) and then toggle the **NoFreeSpaceFlag** back to **OFF**. The following procedure is used to toggle the state of the **NoFreeSpaceFlag**.

Toggle the State of the NoFreeSpaceFlag

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, providing status and capabilities for changing the **Screen Refresh Rate**, **NoFreeSpaceFlag**, and **Suspension** state, as well as a table of active insert processes.
 - 2 To toggle the state of the **NoFreeSpaceFlag** from **ON** to **OFF** (i.e., if the **ON** selection/indicator is filled), click the **OFF** selection/indicator button.
 - The **OFF** selection/indicator button is filled to indicate its selection.
 - 3 Click on the **Apply** button to the right of the **ON** and **OFF** selection/indicator buttons.
 - The screen is refreshed and the **NoFreeSpaceFlag** status line shows "NoFreeSpaceFlag was turned off <date/time>" to indicate the change.
-

Managing Data Pool Configuration Parameters and Data Pool Tuning

The **Manage Configuration Parameters** tab on the DPM GUI permits setting or changing configuration parameters for a number of configuration defaults for the GUI. A prior section of this lesson (Monitoring Data Pool Active Insert Processes and Insert Actions) addressed changing the **Screen Refresh Rate** parameter using an entry field on the **Home Page** tab. This parameter may also be changed using an entry field on the **Manage Configuration Parameters** tab. Although most of the parameters managed on this tab are not likely to be changed frequently, the operator may want to change some of them for tuning the Data Pool. Three of them are addressed in this section. The change process is essentially the same for any of the parameters listed on the tab.

Data Pool tuning parameters can be used to help meter the flow of data into the Data Pool and to adjust retention priority and duration to maintain optimum usage of Data Pool storage. To determine the best settings, it is necessary to monitor Data Pool inserts and disk space and adjust the parameters based on experience and projected functioning.

The first parameter to be considered here is **NumOfAllowedInsertProcesses**, which specifies the maximum number of active insert processes allowed. This number can be an important tuning parameter for regulating the insert-processing stream. As noted earlier, the Data Pool design is based on quick insert processing of data that are available in cache. If too many inserts are placed in the action queue at once, by the time some of those at the end of the queue can be processed, the data may no longer be available in cache and the insert can fail. Use the following procedure to set the number of allowed active insert processes.

Configure the Number of Allowed Active Insert Processes

- 1** Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**).
 - 2** Click on the **Manage Configuration Parameters** tab.
 - The **Manage Configuration Parameters** page is displayed, providing a table of parameters showing three columns: **Parameter Name**, **Parameter Value** (including an entry field with current value, followed by a brief description of the parameter), and **Click on Box to Modify Parm** (containing a check box to mark the parameter for change).
 - There is an **Apply Change** button at the bottom of the page to implement any selected change(s).
 - 3** In the row for the **NumOfAllowedInsertProcesses** parameter, click at the end of the entry field in the **Parameter Value** column.
 - The cursor is displayed at the end of the entry field.
 - 4** Use the **Backspace** key to remove the current value, and then type the desired value.
 - The typed entry is displayed in the field.
 - 5** In the row for the **NumOfAllowedInsertProcesses** parameter, click in the check box in the **Click on Box to Modify Parm** column.
 - The box is filled to indicate selection.
 - 6** Click the **Apply Change** button.
 - The screen is refreshed, the check box is unfilled, and the displayed **Parameter Value** reflects the change.
-

Two other useful Data Pool tuning parameters, addressed together in this section, are **DefaultRetentionPeriod** (for data inserted into the Data Pool based on a Data Pool insert subscription, the default number of days the data are retained in the Data Pool) and **DefaultRetentionPriority** (for data inserted into the Data Pool based on a Data Pool insert subscription, the default retention priority). These parameters are used by the Data Pool Cleanup utility as a basis for retaining data in the Data Pool during cleanup and can be tuned to keep as much data as is practical in the Data Pool without continually running out of space for new inserts. The following procedure may be used to adjust the parameters.

Configure the Default Retention Period and the Default Retention Priority

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**).
- 2 Click on the **Manage Configuration Parameters** tab.
 - The **Manage Configuration Parameters** page is displayed, providing a table of parameters showing three columns: **Parameter Name**, **Parameter Value** (including an entry field with current value, followed by a brief description of the parameter), and **Click on Box to Modify Parm** (containing a check box to mark the parameter for change).
 - There is an **Apply Change** button at the bottom of the page to implement any selected change(s).
- 3 In the row for the **DefaultRetentionPeriod** parameter, click at the end of the entry field in the **Parameter Value** column.
 - The cursor is displayed at the end of the entry field.
- 4 Use the **Backspace** key to remove the current value, and then type the desired value.
 - The typed entry is displayed in the field.
- 5 In the row for the **DefaultRetentionPeriod** parameter, click in the check box in the **Click on Box to Modify Parm** column.
 - The box is filled to indicate selection.
- 6 In the row for the **DefaultRetentionPriority** parameter, click at the end of the entry field in the **Parameter Value** column.
 - The cursor is displayed at the end of the entry field.
- 7 Use the **Backspace** key to remove the current value, and then type the desired value.
 - The typed entry is displayed in the field.

- 8 In the row for the **DefaultRetentionPriority** parameter, click in the check box in the **Click on Box to Modify Parm** column.
 - The box is filled to indicate selection.
 - 9 Click the **Apply Change** button.
 - The screen is refreshed, the check boxes are unfilled, and the displayed **Parameter Value** reflects the changes.
-

Number of Drivers to Run

It is recommended to start with three event drivers, three action drivers, one recovery driver, and one deletion driver, although larger DAACs may eventually need more event drivers and action drivers running. The numbers of drivers are determined by parameters specified in the command to execute the start script for NSBRV drivers. Although one recovery driver and one deletion driver will generally be sufficient, it may be desirable to run more event drivers and action drivers. For example, if increased throughput is needed, it may be appropriate to double the number of action drivers and event drivers to six each. The start script is **EcNbDriverStart**; it is located with other start scripts in the directory **/usr/ecs/<MODE>/CUSTOM/utilities**. To start NSBRV drivers, the command is:

EcNbDriverStart <MODE> d_e d_a d_r d_d

where **<MODE>** is the mode in which Data Pool is being run, **d_e** is the number of event drivers, **d_a** is the number of action drivers, **d_r** is the number of recovery drivers, and **d_d** is the number of deletion drivers. For example, to start the drivers in the OPS mode with the beginning recommendation, use the command **EcNbDriverStart OPS 3 3 1 1**.

To add drivers, the script can simply be run again. For example, to increase the number of event and action drivers to six each, execute the command **EcNbDriverStart OPS 3 3 0 0**.

A database query can give you an idea of the concurrency of processing by event and action drivers -- i.e., how well action processing is keeping up with event processing. You can use *isql* and select **max(actionDateTime** from **EcNbActionQueueLog** where **actionStatus = 'Acquire'** or **actionStatus = 'ActionNotification'**. If the result is more than an hour earlier than the current time, this is an indication that action processing may be lagging behind event processing. In that event, it may be appropriate to try increasing the number of action drivers to one and one-half times the number of event drivers (e.g., six event drivers and nine action drivers). This can result in greater concurrency in processing, but an increase in the number of waiting processes and/or deadlocks occurring on the Sybase server may impose a practical limit on the number of drivers that it is productive to run. Each site will have to monitor Data Pool performance and make adjustments based on experience as necessary to achieve optimum performance.

Managing Data Pool Collection Groups

The conceptual structure of the data pool is set up for each DAAC based on the collections and granules archived at the DAAC. Related collections are grouped into **Collection Groups** (e.g., ASTER collections and granules from the Terra mission, MODIS Oceans collections and granules from the Terra Mission, MISR collections and granules from the Terra mission, MODIS Snow and Ice collections and granules from the Terra mission). Each collection group initially consists of a number of collections that have been specified as valid for Data Pool insertion (i.e., granules of the data types in the collection may be inserted into the Data Pool).

The **Manage Collection Groups** tab of the DPM GUI allows an operator to view and modify collection groups or to add a collection group in the Data Pool database. It also provides access to screens for viewing and modifying collections within a collection group, or for adding a collection to a collection group. Use the following procedure to display the list of collection groups that have collections specified as valid for Data Pool insertion and to view information about those collections.

Use the DPM GUI to View Collection Groups and Collections

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**).
- 2 Click on the **Manage Collection Groups** tab.
 - The **Manage Collection Groups** page is displayed, providing a table listing collection groups with columns providing for each group a **Group ID** and a brief **Description** of the group. The entries in the **Group ID** column are links.
 - At the bottom of the page there are links permitting access to functions for **Add Collection Group** and **Modify Collection Group Description**.

- 3 To obtain more information about the collections in one of the groups, click on its link in the **Group ID** column.
 - The **List of Collections** page is displayed. Depending on whether the collection group is an ECS or a NONECS collection group, the page specifies that the **Data Source** is “ECS” or “NONECS,” respectively. The page identifies the selected group at the top and lists for that group the collections that are valid for insertion in the Data Pool, as well as whether the insertion is for science and metadata or metadata only. On this page, the collection ID entries in the **Collection** column are links.
 - 4 To obtain a description of one of the collections, click on its link in the **Collection** column.
 - A **Description for Collection** page is displayed identifying the selected collection and providing the description. There is a link at the bottom to enable the operator to **Return to previous page**.
 - 5 If it is desirable to view a description for another collection in the same group, click on the **Return to previous page** link; otherwise, go to Step 7.
 - The **List of Collections** page is displayed again.
 - 6 Repeat Steps 4 and 5 as desired.
 - 7 If it is desired to obtain more information about collections in a different collection group, repeat Steps 2 through 6.
-

Rarely, it may be desirable to modify the description of one or more of the collection groups listed on the **Manage Collection Groups** page. If there is a need to modify a collection group description, there is a link at the bottom of the list on that page providing access to a page that permits the descriptions to be modified. The following procedure is applicable.

Use the DPM GUI to Modify Collection Groups

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**).
- 2 Click on the **Manage Collection Groups** tab.
 - The **Manage Collection Groups** page is displayed, providing a table listing collection groups with columns providing for each group a **Group ID** and a brief **Description** of the group. The entries in the **Group ID** column are links.
 - At the bottom of the page there are links permitting access to functions for **Add Collection Group** and **Modify Collection Group Description**.

- 3 Click on the **Modify Collection Group Description** link at the bottom of the page.
 - The screen displays a page listing the collection groups with their descriptions in text-entry fields, each accompanied by a check box in a **Click on Box to Modify Desc** column.
 - 4 Click at the end of the text-entry field for the collection group to be modified.
 - The cursor is displayed at the end of the **Description** text-entry field.
 - 5 Use the **Backspace** key to delete the existing description and type the desired description (*Note:* You may also use the mouse to drag the cursor over all or part of the text to be replaced, highlighting the selected text, and type new text to replace the highlighted text.)
 - The typed text appears in the field.
 - 6 Click on the check box in the **Click on Box to Modify Desc** column to the right of the newly modified description.
 - The box is filled to indicate its selection.
 - 7 Repeat Steps 4 through 6 for any additional descriptions to be modified.
 - 8 Click on the **Apply Change** button at the bottom of the **Click on Box to Modify Desc** column.
 - The screen is refreshed and the **Manage Collection Groups** page reflects the changed description(s).
-

From time to time, it may be necessary to add a collection group (e.g., if a DAAC begins archiving data from a new instrument). If a collection group is to be added to the list of collection groups, it is necessary to use the **Add Collection Group** link at the bottom of the **Manage Collection Groups** page. *NOTE:* Although the following procedure is applicable, most of the time new collection groups will be added only during releases of new software versions and you will not typically use this procedure.

Caution:

The Add Collection Group function is to be exercised judiciously because the DPM GUI does not provide any means of deleting collection groups.

Use the DPM GUI to Add a Collection Group

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).

- The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**).
- 2 Click on the **Manage Collection Groups** tab.
 - The **Manage Collection Groups** page is displayed, providing a table listing collection groups with columns providing for each group a **Group ID** and a brief **Description** of the group. The entries in the **Group ID** column are links.
 - At the bottom of the page there are links permitting access to functions for **Add Collection Group** and **Modify Collection Group Description**.
 - 3 Click on the **Add Collection Group** link at the bottom of the page.
 - The screen displays a page with columns of text-entry fields; a **Group ID** column and a **Description** column permit identification and description of new collection groups.
 - 4 Click on the option button in the **Data Source** column and then click on **ECS** or **NONECS** to specify whether the collection to be added is an ECS Collection or a NONECS Collection, respectively.
 - The selected option is displayed on the option button.
 - 5 Click in a **Group ID** text-entry field.
 - The cursor is displayed in the **Group ID** text-entry field.
 - 6 Type the identifier for a new collection group.
 - The typed entry is displayed in the field.
 - 7 Click in the **Description** text-entry field for the newly entered **Group ID**.
 - The cursor is displayed at the end of the **Description** text-entry field.
 - 8 Type the description for the new collection group.
 - The typed entry is displayed in the field.
 - 9 Repeat Steps 4 through 8 for any additional collection groups to be added.
 - 10 Click on the **Apply Change** button at the bottom of the entry-field area.
 - The screen is refreshed and the **Manage Collection Groups** page reflects the new collection group(s).
-

Managing Data Pool Collections within Collection Groups

Although an initial Data Pool structure is provided, not all collections are necessarily specified as eligible for Data Pool insertion. Based on experience, or on changes in demand, a DAAC

may wish to add one or more collections to a data group. Use the following procedure to add collections to an existing ECS collection group.

Use the DPM GUI to Add ECS Collections to a Collection Group

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**).
- 2 Click on the **Manage Collection Groups** tab.
 - The **Manage Collection Groups** page is displayed, providing a table listing collection groups with columns providing for each group a **Group ID** and a brief **Description** of the group. The entries in the **Group ID** column are links.
- 3 Click on the **Group ID** link for the ECS collection group to which the collection is to be added.
 - The **List of Collections** page is displayed, identifying the selected group at the top and listing for that group the collections and versions that are valid for insertion in the Data Pool, as well as whether the insertion is for science and metadata or metadata only. The page also lists whether the collections are enabled for HEG processing, and the Spatial Search Type for each collection.
- 4 Click on the **Add Collection** link at the bottom.
 - The **Add Collections** page is displayed, indicating the selected collection group at the top and listing from the Science Data Server database collections not currently approved as eligible for Data Pool insertion. The page also provides a column with option buttons permitting the operator to select **Science and Metadata** or **Metadata Only** and a column enabling selection of **Valid for data pool** or **Invalid for data pool**. The defaults are **Science and Metadata** and **Invalid for data pool**, respectively. There is also a **Click on Box to Add Collection** column with check boxes to mark the collection(s) for addition, and an **Apply Change** button to implement the addition(s).
- 5 If the collection to be added is to permit insertion of **Science and Metadata**, go to Step 6; otherwise, click on its option button in the **Science Granules and/or Metadata** column and click to select **Metadata Only**.
 - The selected choice is displayed on the option button.
- 6 Click on the collection's option button in the **Data Pool Insertion** column and click to select **Valid for data pool**.
 - The selected choice is displayed on the option button.

- 7 Click on the collection's check box in the **Click on Box to Add Collection** column.
 - The box is filled to indicate its selection.
 - 8 Repeat Steps 5 through 7 for any additional collections to be added to the collection group identified at the top of the page.
 - 9 Click on the **Apply Change** button at the bottom of the **Click on Box to Add Collection** column.
 - The **List of Collections** page reflects the addition(s).
 - 10 Repeat Steps 2 through 9 for any additional collection(s) to be added to another collection group.
-

Use the DPM GUI to Add NONECS Collections to a Collection Group

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**).
- 2 Click on the **Manage Collection Groups** tab.
 - The **Manage Collection Groups** page is displayed, providing a table listing collection groups with columns providing for each group a **Group ID** and a brief **Description** of the group. The entries in the **Group ID** column are links.
- 3 Click on the **Group ID** link for the NONECS collection group to which the collection is to be added.
 - The **List of Collections** page is displayed, identifying the selected group at the top and listing for that group the collections and versions that are valid for insertion in the Data Pool, as well as whether the insertion is for science and metadata or metadata only. The page also lists whether the collections are enabled for HEG processing, and the Spatial Search Type for each collection.

- 4 Click on the **Add Collection** link at the bottom.
 - The **Add Collections** page is displayed, indicating the selected collection group at the top and permitting entry or specification of information to define the collection to be added. The page provides a column with an entry field for **Collection**, a column with an entry field for **Version**, and a column with an entry field for **Description**. It provides a column with an option button permitting the operator to select **Science and Metadata** or **Metadata Only** and a column enabling selection of **Valid for data pool** or **Invalid for data pool**. The defaults are **Science and Metadata** and **Valid for data pool**, respectively. There is also a column with an option button permitting selection of a **Spatial Search Type**; the default is **Not supported**. An **Apply Change** button is provided at the bottom to implement the addition.
 - 5 Click in the entry field for **Collection**.
 - The cursor is displayed in the field.
 - 6 Type the name for the collection.
 - The typed entry is displayed in the field.
 - 7 Click in the entry field for **Version**.
 - The cursor is displayed in the field.
 - 8 Type the version.
 - The typed entry is displayed in the field.
 - 9 If the collection to be added is to permit insertion of **Science and Metadata**, go to Step 10; otherwise, click on its option button in the **Science Granules and/or Metadata** column and click to select **Metadata Only**.
 - The selected choice is displayed on the option button.
 - 10 Click on the option button in the **Data Pool Insertion** column and click to select **Valid for data pool**.
 - The selected choice is displayed on the option button.
 - 11 Click on the option button in the **Spatial Search Type** column and click to select the spatial search type (choices are **Not supported**, **Gpolygon**, **Rectangle**, and **Orbit**).
 - The selected choice is displayed on the option button.
 - 12 Click on the **Apply Change** button.
 - The **List of Collections** page reflects the addition(s).
-

As part of managing the Data Pool storage and retention of data, making adjustments based on experience and/or changes in demand, it may be desirable to modify a collection. The modification may mean specifying that metadata only may continue to be inserted and science granules may no longer be inserted, or declaring the collection no longer valid for data pool insertion at all. Use the following procedure for modifying a collection.

Use the DPM GUI to Modify Collections in a Collection Group

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**).
- 2 Click on the **Manage Collection Groups** tab.
 - The **Manage Collection Groups** page is displayed, providing a table listing collection groups with columns providing for each group a **Group ID** and a brief **Description** of the group. The entries in the **Group ID** column are links.
- 3 Click on the **Group ID** link for the collection group containing the collection to be modified.
 - The **List of Collections** page is displayed, identifying the selected group at the top and listing for that group the collections and versions that are valid for insertion in the Data Pool, as well as whether the insertion is for science and metadata or metadata only. The page also lists whether the collections are enabled for HEG processing, and the Spatial Search Type for each collection.
- 4 Click on the **Modify Collection** link at the bottom.
 - The **Modify Collections** page is displayed, indicating the selected collection group at the top, its data source (ECS or NONECS) and listing collections currently approved as eligible for Data Pool insertion. The page also provides a column with option buttons permitting the operator to select **Science and Metadata** or **Metadata Only** and a column enabling selection of **Valid for data pool** or **Invalid for data pool**. If the data source is NONECS, there is also a column with option buttons permitting selection of **Spatial Search Type**.
- 5 If the modification to a collection is to change the selected option displayed on the option button for insertion of **Science and Metadata** (in the **Science Granules and/or Metadata** column), click on that option button for the collection and click to select the desired option from the displayed choices.
 - The selected choice is displayed on the option button.

- 6 If the modification is to change a collection's eligibility for insertion in the Data Pool, click on that collection's option button in the **Data Pool Insertion** column and click to select the desired option from the displayed choices.
 - The selected choice is displayed on the option button.
 - 7 If the modification is for a NONECS Collection and is to change a collection's Spatial Search Type, click on that collection's option button in the **Spatial Search Type** column and click to select the desired option from the displayed choices.
 - The selected choice is displayed on the option button.
 - 8 Click on the collection's check box in the **Click on Box to Modify Collection** column.
 - The box is filled to indicate its selection.
 - 9 Repeat Steps 5, 6, and 8 for any additional ECS collections to be modified or Steps 5 through 8 for any additional NONECS collections to be modified.
 - 10 Click on the **Apply Change** button at the bottom of the **Click on Box to Modify Collection** column.
 - The **List of Collections** page reflects the change(s).
 - 11 Repeat Steps 2 through 10 for any modifications to one or more collections in another collection group.
-

Using the DPM GUI to Manage Themes

The Synergy III release provides the capability for users to search the Data Pool for data associated with themes. Accordingly, as data are inserted into the Data Pool, it is possible to have the data associated with themes. The DPM GUI Manage Themes tab permits viewing, creating, modifying, and deleting themes. The following procedures are applicable.

Use the DPM GUI to View a List of Themes

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**).
- 2 Click on the **Manage Themes** tab.
 - The **Manage Themes** page is displayed, providing a table listing themes with columns providing for each theme a **Theme Name** and **Description**, an indication of whether the theme is **Insert Enabled**, an indication of whether the theme is **Web Enabled**, and a check box that can be used with an **Apply Change** button at the

- bottom to delete the theme. There are option buttons, a text entry field, and an **Apply Filter** button at the top for specifying filters to limit the number of themes displayed, permitting selection on the basis of whether the theme is **Web Visible** or **Insert Enabled**, or by the **Beginning Letters** in the theme name.
- 3 If it is desirable to filter the list, click either or both option buttons to select the desired **Web Visible** and/or **Insert Enabled** option and/or click in the **Beginning Letters** field and type the first few letters of the name(s) of known and sought themes.
 - The option button(s) display the selected choice(s) and/or the **Beginning Letters** field displays the typed entry.
 - 4 Click on the **Apply Filters** button.
 - The listed themes reflect the choices displayed for the filters.
-

The DPM GUI may be used to modify a theme. This can be useful if, for example, it is noted that access frequency for granules referencing a theme has declined to the point that the thematic collection should be removed from the Data Pool, but there are still a few web users that still use it. In that case, it may be appropriate to change the description of the theme to alert users that the theme will soon be phased out. Use the following procedure.

Use the DPM GUI to Modify a Theme

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**).
- 2 Click on the **Manage Themes** tab.
 - The **Manage Themes** page is displayed, providing a table listing themes with columns providing for each theme a **Theme Name** and **Description**, an indication of whether the theme is **Insert Enabled**, an indication of whether the theme is **Web Enabled**, and a check box that can be used with an **Apply Change** button at the bottom to delete the theme. There are option buttons, a text entry field, and an **Apply Filter** button at the top for specifying filters to limit the number of themes displayed, permitting selection on the basis of whether the theme is **Web Visible** or **Insert Enabled**, or by the **Beginning Letters** in the theme name.
- 3 Click on the **Modify Theme** link at the bottom of the list of themes (scrolling down if necessary).
 - The **Modify Theme** page is displayed.

- 4 If it is desirable to filter the displayed list of themes and characteristics, click either or both option buttons to select the desired **Web Visible** and/or **Insert Enabled** option and/or click in the **Beginning Letters** field and type the first few letters of the name(s) of known and sought themes.
 - The option button(s) display the selected choice(s) and/or the **Beginning Letters** field displays the typed entry.
 - 5 Click on the **Apply Filters** button
 - The listed themes reflect the choices displayed for the filters.
 - 6 If it is desirable to modify the description of a theme, click in the field in the **Description** column in the row for the theme.
 - The cursor is displayed in the field.
 - 7 Modify the text as desired, either by typing over text highlighted by dragging the cursor or by backspacing to delete undesired text and typing the desired replacement.
 - The desired description is displayed in the field.
 - 8 If it is desired to change whether the theme is **Web Visible** and/or **Insert Enabled**, click on the toggle button box in the appropriate column(s) in the row for the theme
 - Any filled button boxes that are clicked toggle to unfilled, and any unfilled button boxes that are clicked toggle to filled.
 - 9 Click on the toggle button box in the **Click on Box to Modify** column in the row for the theme being changed.
 - The clicked toggle button box is filled.
 - 10 Repeat steps 4 –9 for any additional themes to be modified.
 - 11 Click on the **Apply Change** button at the bottom of the page.
 - Any changes are applied in the database and the box(es) in the **Click on Box to Modify** column that were filled are unfilled.
-

Use the following procedure to add a theme.

Use the DPM GUI to Add a Theme

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**).

- 2 Click on the **Manage Themes** tab.
 - The **Manage Themes** page is displayed, providing a table listing themes with columns providing for each theme a **Theme Name** and **Description**, an indication of whether the theme is **Insert Enabled**, an indication of whether the theme is **Web Enabled**, and a check box that can be used with an **Apply Change** button at the bottom to delete the theme. There are option buttons, a text entry field, and an **Apply Filter** button at the top for specifying filters to limit the number of themes displayed, permitting selection on the basis of whether the theme is **Web Visible** or **Insert Enabled**, or by the **Beginning Letters** in the theme name.
 - 3 Click on the **Add New Theme** link at the bottom of the list of themes (scrolling down if necessary).
 - The **Add New Theme** page is displayed.
 - 4 Click in the **Theme** field in the **Theme Name** column.
 - The cursor is displayed in the field.
 - 5 Type a name for the theme to be added.
 - The typed entry is displayed in the field.
 - 6 Click in the **Description** field in the **Description** column.
 - The cursor is displayed in the field.
 - 7 Type a description for the theme to be added.
 - The typed entry is displayed in the field.
 - 8 If it is desirable to make the theme web visible at the current time, click in the toggle button box in the **Web Visible** column.
 - The clicked box is filled.
 - 9 If it is desirable to enable Data Pool inserts for the theme at the current time, click in the toggle button box in the **Insert Enabled** column.
 - The clicked box is filled.
 - 10 Click on the **Apply Change** button at the bottom of the page.
 - The theme is added to the database and the entry fields are cleared.
 - 11 Click on the **Return to theme list** link.
 - The added theme is displayed in the list of themes.
-

Use the following procedure to delete a theme.

Use the DPM GUI to Delete a Theme

- 1 Launch the DPM GUI (refer to procedure **Launch the DPM GUI**, previous section of this lesson).
 - The **Home Page** is the default display, offering tabs for access to Data Pool maintenance functions (**Batch Summary**, **List Insert Queue**, **Manage Configuration Parameters**, **Manage Collection Groups**, and **Manage Themes**).
 - 2 Click on the **Manage Themes** tab.
 - The **Manage Themes** page is displayed, providing a table listing themes with columns providing for each theme a **Theme Name** and **Description**, an indication of whether the theme is **Insert Enabled**, an indication of whether the theme is **Web Enabled**, and a check box that can be used with an **Apply Change** button at the bottom to delete the theme. There are option buttons, a text entry field, and an **Apply Filter** button at the top for specifying filters to limit the number of themes displayed, permitting selection on the basis of whether the theme is **Web Visible** or **Insert Enabled**, or by the **Beginning Letters** in the theme name.
 - 3 If it is desirable to filter the list, click either or both option buttons to select the desired **Web Visible** and/or **Insert Enabled** option and/or click in the **Beginning Letters** field and type the first few letters of the name(s) of known and sought themes.
 - The option button(s) display the selected choice(s) and/or the **Beginning Letters** field displays the typed entry.
 - 4 Click on the **Apply Filters** button.
 - The listed themes reflect the choices displayed for the filters.
 - 5 In the row for the theme to be deleted, click in the toggle button box in the **Click on Box to Delete** column.
 - The clicked box is filled.
 - 6 Repeat step 5 for any additional themes to be deleted at this time.
 - 7 Click on the **Apply Change** button at the bottom of the page.
 - A dialog box asks **Are you sure you want to delete this theme?**
 - 8 Click on the **OK** button in the dialog box.
 - The list of themes reflects the deletion(s).
-

Troubleshooting DPM GUI Problems

The Data Pool maintenance modules use Perl scripts and Sybase database functions. If the tool cannot be launched, or does not function (e.g., cannot manage collection groups, does not update screens), you will need to ask the System Administrator to ensure that the appropriate scripts are functioning properly. It may be necessary to have the Database Administrator check to ensure that there are no problems with the database.

It is also possible to receive error messages when using the DPM GUI while it is apparently functioning normally. Error messages associated with the DPM GUI are listed in Appendix A of the *Operations Tools Manual* (Document 609-CD-610-003, Table A.2.11-1). Table 5 is taken from the corresponding table in Document 609. If a problem cannot be identified and fixed without help within a reasonable period of time, the appropriate response is to call the help desk or submit a trouble ticket in accordance with site Problem Management policy.

Table 5. DPM GUI User Messages

Message Text	Impact	Cause and Corrective Action
DB Error: You entered a duplicate collection group name that exists in the database. Please try again	Unable to add a new group id	Duplicate group name is entered. Check the list of group ids and enter a group name consisting of four letters, which is not on the list.
INPUT Error: You entered an invalid group name. Please see help page for more information. Please see section Add Collection Group	Unable to add a new group	Lower case letter is entered. Group id should be all Upper case letters.
INPUT Error: You entered an invalid name. Please see log for more details. Consult help tab and see section for NON-ECS add Collection Screen	Unable to add a new non-ECS collection	Special characters/small letters are entered. Non-ECS collection name should be in capital letters and without any special characters.
INPUT Error: You entered an invalid theme name. Please see help page for more information. Please see section: Add New Theme	Unable to add a new theme	Special characters/small letters are entered. Theme names should be in capital letters and without any special characters.
DB Error: Theme can not be null or empty	Unable to add a theme	A null or empty string is entered. Theme name should contain capital, small letters. Space is also allowed but no special characters.

Table 5. DPM GUI User Messages (Cont.)

Message Text	Impact	Cause and Corrective Action
DB Error: You entered either an existing theme name or a collection or a group name or an ESDT name. Check the log at /usr/ecs/<mode>/CUSTOM/log/EcDIDpmDataPoolGui.log for more details	Unable to add a theme	A name is entered, which is a duplicate name for a group, collection or an ESDT name.
DB Error: This collection is allowed for insertion therefore Spatial Search Type cannot be modified for this collection	Unable to modify Spatial Search type for a collection	Collection is not allowed for insertion. First make the collection allowed for insertion and then try to modify search type
DB Error: Error adding this collection. Collection entry <collection name> <version> already exist	Unable to add a collection.	Duplicate collection name entered. Verify the list of collection and then enter a name, which is unique.
DB Error: Internal error occurred	A db transaction interrupted.	Database connection is lost for network error. No suggestion.
DB Error: delete failed because there are granules associated with this theme	Unable to delete a theme.	There are granules associated with this theme. Disassociate granules from this theme and then delete it.

Using the Spatial Subscription Server (NSBRV) GUI

Although there are some Data Pool maintenance tasks that are done with scripts and utilities, as described in a subsequent section of this lesson, the Spatial Subscription Server (NSBRV) GUI provides a convenient means of modifying the period of retention in a Data Pool insert subscription. This can be done by User Services and/or science personnel, but the archive support personnel should also be familiar with the GUI and its use. Other tasks done with the NSBRV GUI may also be of use to archive personnel (e.g., viewing the acquire and notification actions currently being processed by the NSBRV, viewing statistics on the processing of events and actions by the NSBRV).

The NSBRV GUI is a web application; Figure 48 illustrates the NSBRV GUI **Home Page**, from which the operator can navigate to other pages to begin specific tasks. As the figure shows, in addition to the **Home Page** tab, it provides four tabs for access to pages supporting various tasks:

- **List Events:** access to pages for listing subscribable events.
- **Manage Subscriptions:** access to pages for managing subscriptions.
- **Manage Bundling Orders:** access to pages for adding, viewing, updating and listing subscriptions for bundling orders.
- **Monitor Queues:** access to pages for monitoring runtime queues and production statistics.

There is also a Help tab providing descriptions of the NSBRV functions to provide the operator with assistance in navigating through the GUI.

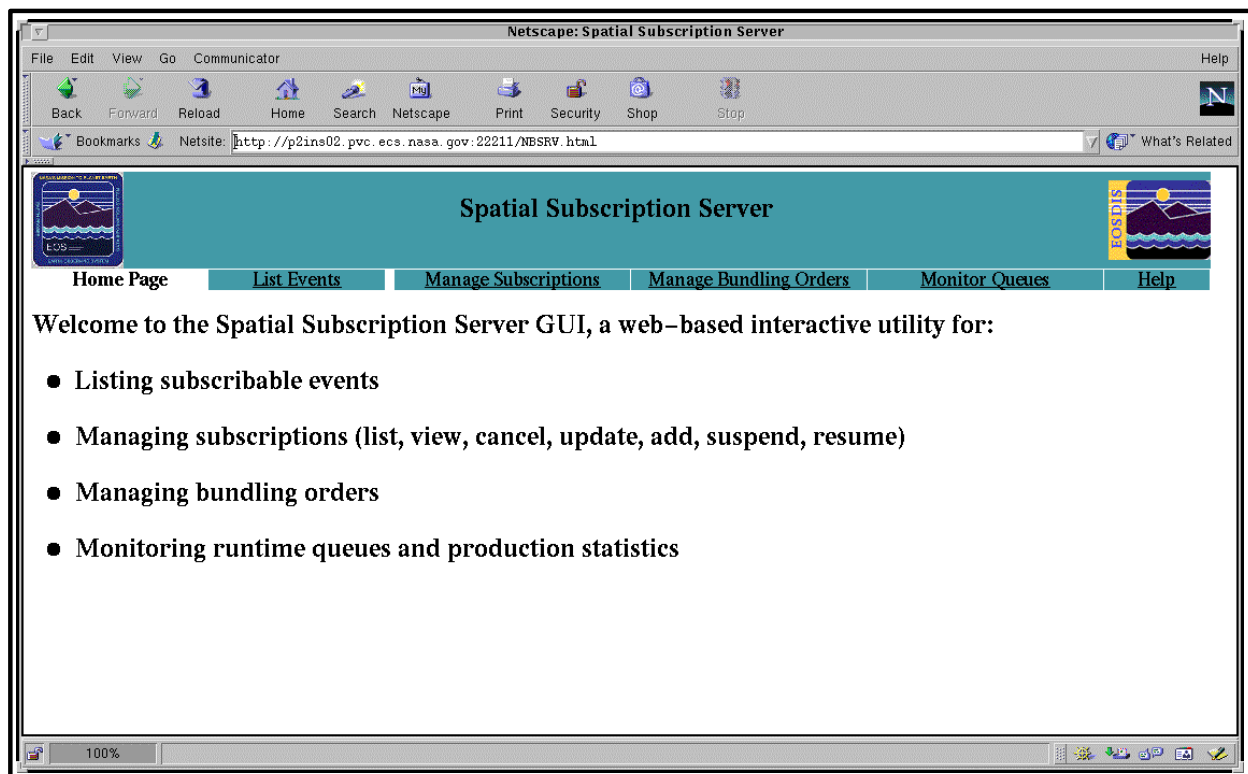


Figure 48. NSBRV Home Page

Figure 49 shows the **Manage Subscriptions** page. This page lists subscriptions existing in the NSBRV database, with columns for **Subscription Id**, **User** for whom the subscription was created, **Collection**, **Version**, **Event Type**, **Status**, and **Expiration Date**. These column headers, except for **Version** and **Event Type**, are links on which the list can be sorted. There are also three option buttons and a **Filter** button, permitting the operator to filter the list by user, collection, and status. Finally, at the right side of the page, there is a **Choose Subscription Action** column permitting operators an option, for each listed subscription, to **View**, **Update**, or **Delete** the subscription, with an **Apply** button to implement a selected option. Note also that two additional tabs appear on this page. The **Add Subscriptions** tab is employed by data specialists/User Services or other personnel to add new subscriptions

Spatial Subscription Server

Home Page | List Events | **Manage Subscriptions** | Add Subscriptions | List Themes | Monitor Queues | Help

User: Collection: Status:

Subscription Id	User	Collection	Version	Event Type	Status	Expiration Date	Choose Subscription Action
3	cmops	AIRH0ScE	001	INSERT	Inactive	Nov 6 2002 12:00AM	<input type="button" value="View"/> <input type="button" value="Update"/> <input type="button" value="Delete"/> <input type="button" value="Apply"/>
5	LarcIngMgr	MOD03	001	INSERT	Active	Nov 7 2002 12:00AM	<input type="button" value="View"/> <input type="button" value="Update"/> <input type="button" value="Delete"/> <input type="button" value="Apply"/>
7	LarcIngMgr	MOD04_L2	001	INSERT	Active	Nov 7 2002 12:00AM	<input type="button" value="View"/> <input type="button" value="Update"/> <input type="button" value="Delete"/> <input type="button" value="Apply"/>
10	NOAA/SOAP	AM1EPHN0	001	INSERT	Active	Nov 7 2002 12:00AM	<input type="button" value="View"/> <input type="button" value="Update"/> <input type="button" value="Delete"/> <input type="button" value="Apply"/>
11	NOAA/SOAP	GDAS_0ZF	001	INSERT	Active	Nov 7 2002 12:00AM	<input type="button" value="View"/> <input type="button" value="Update"/> <input type="button" value="Delete"/> <input type="button" value="Apply"/>
12	NOAA/SOAP	OZONEEP	001	INSERT	Active	Nov 8 2002 12:00AM	<input type="button" value="View"/> <input type="button" value="Update"/> <input type="button" value="Delete"/> <input type="button" value="Apply"/>
14	NOAA/SOAP	SEA_ICE	001	INSERT	Active	Nov 8 2002 12:00AM	<input type="button" value="View"/> <input type="button" value="Update"/> <input type="button" value="Delete"/> <input type="button" value="Apply"/>
15	cgili	DFLAXMIS	001	INSERT	Active	Nov 8 2002 12:00AM	<input type="button" value="View"/> <input type="button" value="Update"/> <input type="button" value="Delete"/> <input type="button" value="Apply"/>

Figure 49. NSBRV Manage Subscriptions page

If an operator selects and implements the **View** option for a subscription, the GUI provides a display similar to that illustrated in Figure 50. The page lists detailed information about the selected subscription.

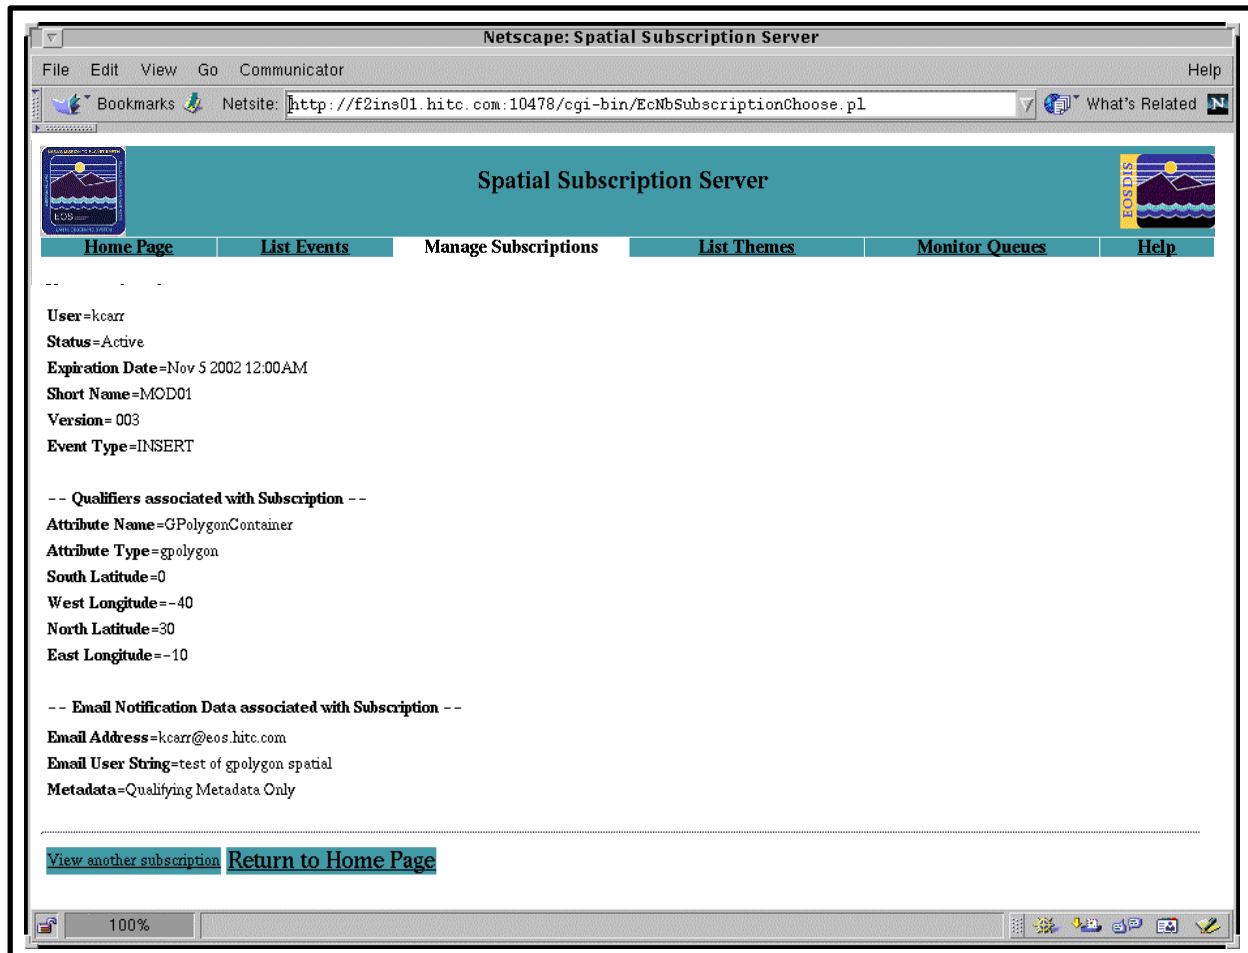


Figure 50. NSBRV View Subscriptions page

If the operator clicks on the **List Themes** tab on the **Manage Subscriptions** page or when viewing one of the subscriptions (i.e., from the **View Subscription** page), the GUI displays the page illustrated in Figure 51, showing a list of themes with which subscriptions may be associated and permitting the operator to list the subscriptions currently associated with any available theme.

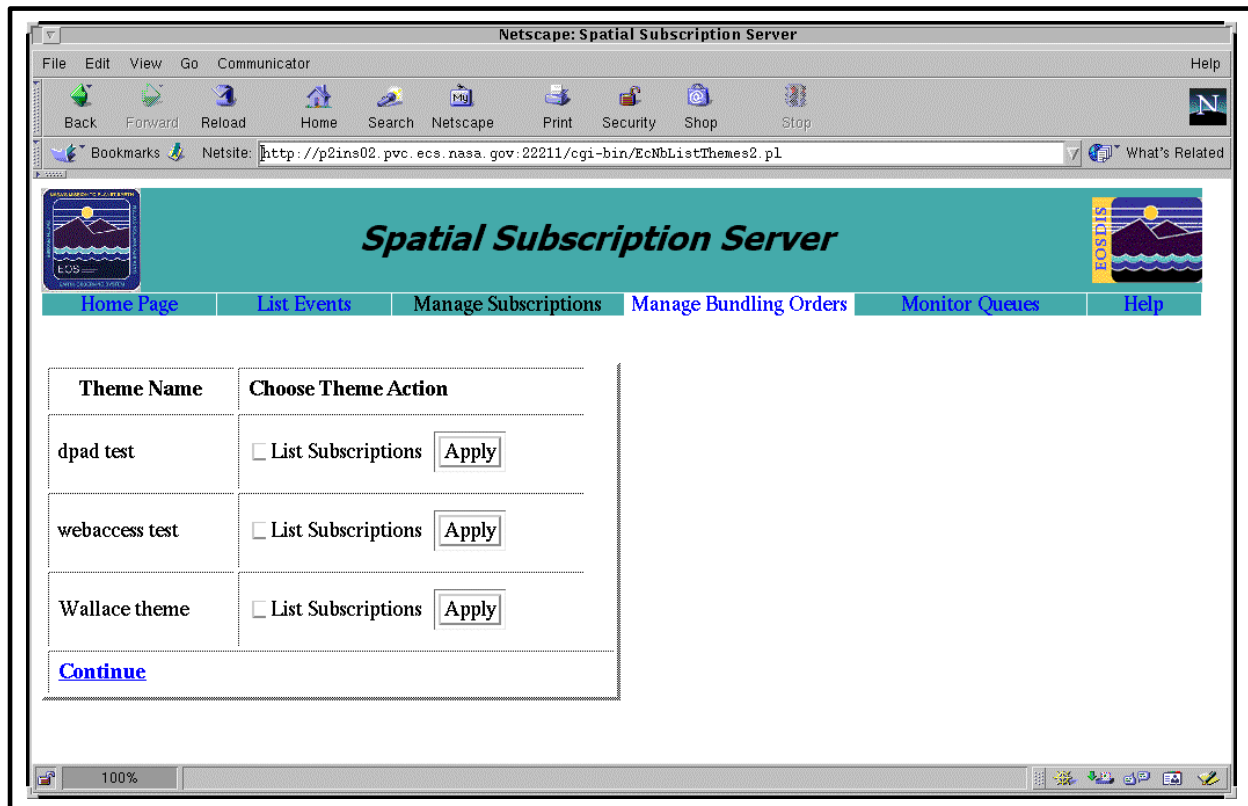


Figure 51. NSBRV List Themes page

If the operator clicks on the **Manage Bundling Orders** tab on the **Home Page** or other page where the tab is available, the GUI displays the page illustrated in Figure 52, listing bundling orders. The list includes a **Bundling Order** identifying number for each bundling order, the **User** for whom the bundling order was created, the **Creation Date** and **Expiration Date**, the **Media Type**, and **Status**. There is an **Add Bundling Order** link to a page where the operator can add a bundling order and a **Configure Defaults** link to a page where the operator can set default parameters for configuring bundling orders. There are also three option buttons and a **Filter** button, permitting the operator to filter the list by user, media type, and status. Finally, at the right side of the page, there is a **Choose Bundling Order Action** column permitting operators an option, for each listed bundling order, to **View**, **Update**, **Cancel**, or **List Subs** (list the subscriptions) associated with the bundling order, with an **Apply** button to implement a selected option. The figure shows the window that is obtained by selecting and implementing the **View** option.

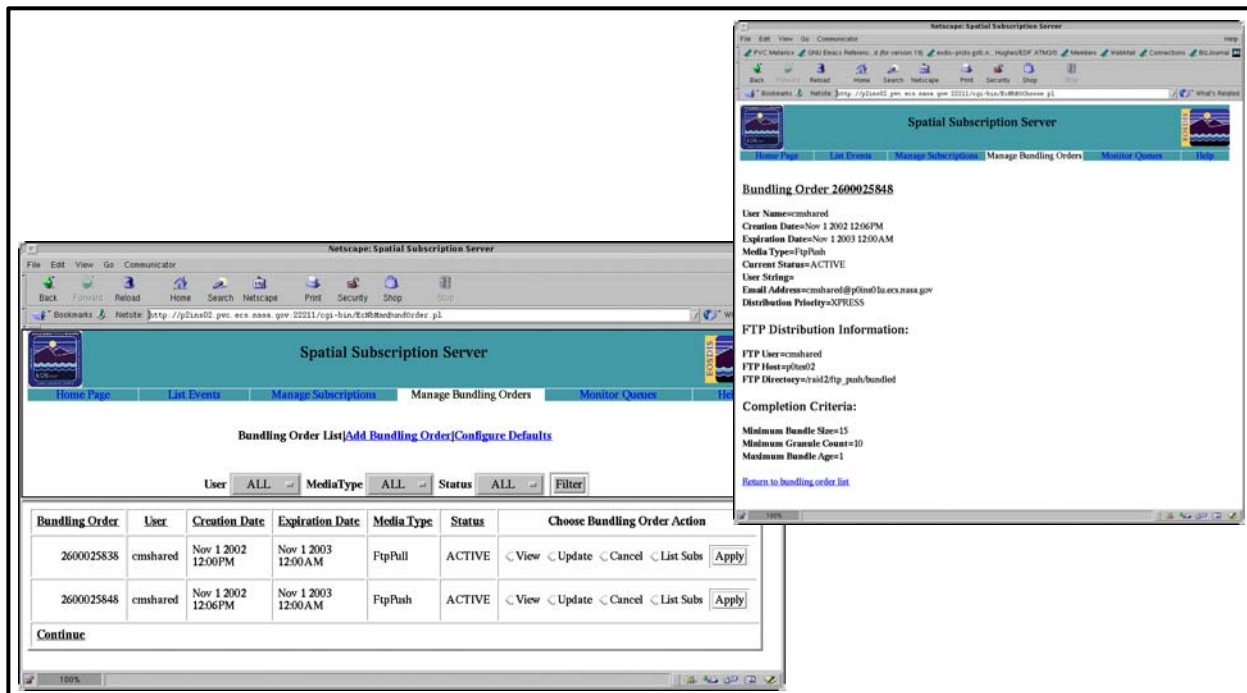


Figure 52. NSBRV Manage Bundling Orders page, with View Window

If the operator clicks on the **Add Bundling Order** link, the GUI displays the page illustrated in Figure 53. This page contains a field for entry of the operator's **UserID**, a field for specification of an **Expiration Date** for the bundling order (with a default entry), and an option button for selection of a media type (CD ROM, 8 mm tape, DVD, DLT, FtpPush or FtpPull). There is a **continue** button at the bottom to continue to a page for further specification of data related to the bundling order; the data to be entered on the resulting page varies with the media type selection. There is also a **clear** button that can be used to clear the **UserID** field and restore the default **Expiration Date** and **Media Type**.

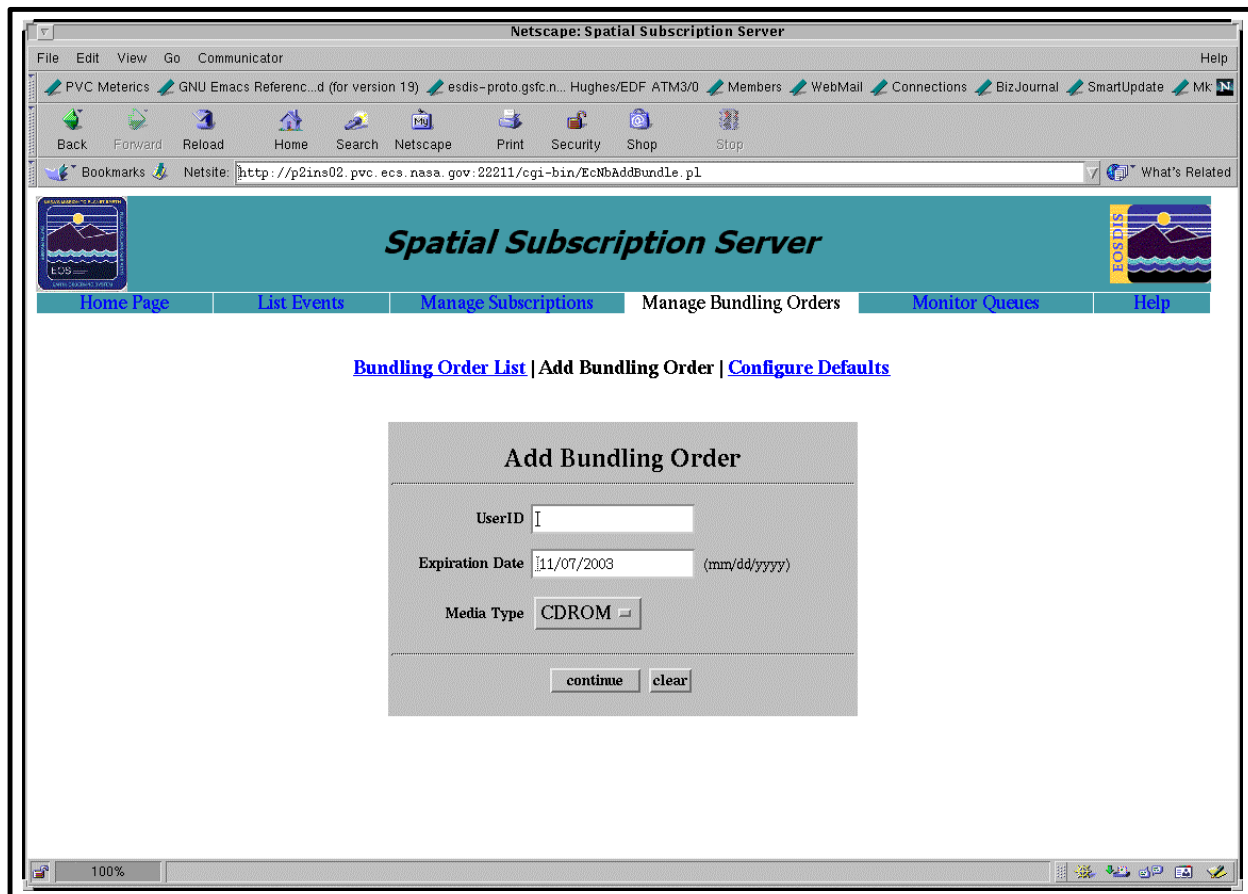


Figure 53. NSBRV GUI Add Bundling Order page

Figure 54 shows the data pages that result from clicking on the **continue** button at the **Add Bundling Order** page when the **Media Type** chosen is a hard media type, an FTP Push, or an FTP Pull, respectively. These pages permit entry of information necessary for full specification of the selected **Media Type**. Each one also requires specification of three criteria for completion of a bundle, whereby fulfillment of any one of the three criteria will trigger distribution of the bundle. The criteria include:

- Minimum bundle size (GB).
- Minimum Granule Count.
- Maximum Bundle Age (days).

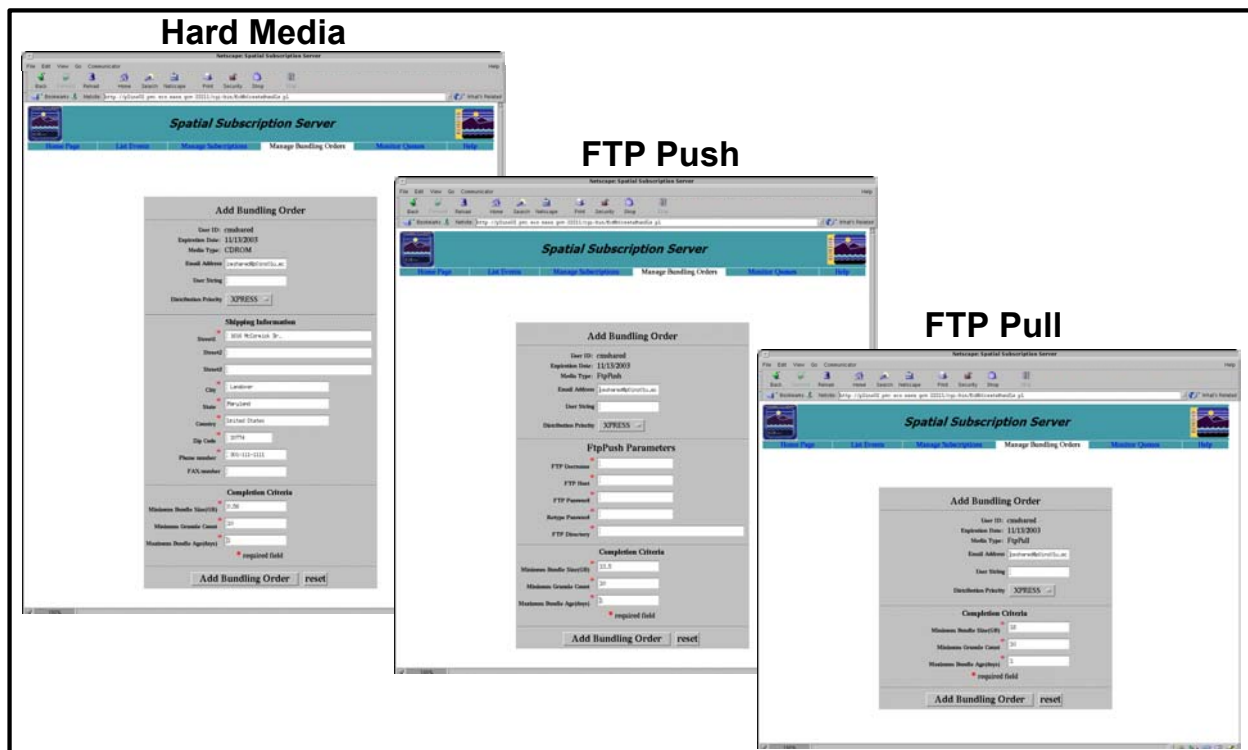


Figure 54. NSBRV GUI Data Pages for adding bundling orders

The GUI provides entry screens for an operator to add a subscription (accessible through the **Add Subscription** tab on the **Manage Subscriptions** page); this task is typically performed by User Services or Science Data personnel and therefore the **Add Subscription** form is not illustrated here. However, if an operator selects and implements the **Update** option for a subscription (listed on the **Manage Subscriptions** page), the GUI provides a display similar to that illustrated in Figure 55. This page is structured very similarly to the **Add Subscription** page, permitting the operator to add or change qualifiers and change the action selection for the subscription. It also permits the operator to specify a bundling order for the subscription. Changes are submitted by clicking on the **Update Subscription** button at the bottom of the page.

Otherwise choose one or more actions ☐ Acquire ☐ E-Mail Notification ☐ Data Pool
Enter information below for the actions selected (not applicable for bundled subscriptions)

Acquire Information

User Profile:

User String:

Email Address:

Media Format: FILEFORMAT

Media Type:

Priority:

Notify Type: MAIL

FTP Information (for FtpPush Only)

FTP User:

FTP Password:

Enter password again for verification:

FTP Host:

FTP Directory:

E-Mail Notification Information

Action Address:

User String:

Metadata:

Data Pool Information

Retention Period: (in days)

Retention Priority: (valid range 1 thru 255)

Science Granules and/or Metadata:

Check here to add theme: ☐ Enter first few chars of name:

Update Subscription

Figure 55. NSBRV Update Subscriptions page

There are two other major pages of the NSBRV GUI accessible from tabs on the **Home Page**. The **List Action Queue** page, illustrated in Figure 56, provides a table listing acquire and notification actions that are being processed. On this page, the **Action Type** and **Subscription Id** column headers are links for sorting the list, and there are also **Action Type**, **Subscription**, and **Status** option buttons and a filter button for filtering the list.

Action Type	Subscription Id	User	Collection	Version	Enqueue Time	Dequeue Time	Status
Notify	13	kcarr	MOD04_L2	003	Dec 31 2001 11:42AM	Dec 31 2001 11:42AM	FinishedActionProcessing
Notify	13	kcarr	MOD04_L2	003	Dec 31 2001 11:42AM	Dec 31 2001 11:42AM	FinishedActionProcessing
Notify	20	kcarr	MOD05_L2	003	Dec 31 2001 11:42AM	Dec 31 2001 11:42AM	FinishedActionProcessing
Notify	20	kcarr	MOD05_L2	003	Dec 31 2001 11:42AM	Dec 31 2001 11:42AM	FinishedActionProcessing
Notify	20	kcarr	MOD05_L2	003	Dec 31 2001 11:42AM	Dec 31 2001 11:42AM	FinishedActionProcessing
Notify	20	kcarr	MOD05_L2	003	Dec 31 2001 11:42AM	Dec 31 2001 11:42AM	FinishedActionProcessing
Notify	20	kcarr	MOD05_L2	003	Dec 31 2001 11:42AM	Dec 31 2001 11:42AM	FinishedActionProcessing
Notify	20	kcarr	MOD05_L2	003	Dec 31 2001 11:42AM	Dec 31 2001 11:42AM	FinishedActionProcessing
Notify	20	kcarr	MOD05_L2	003	Dec 31 2001 11:42AM	Dec 31 2001 11:42AM	FinishedActionProcessing
Notify	20	kcarr	MOD05_L2	003	Dec 31 2001 11:42AM	Dec 31 2001 11:42AM	FinishedActionProcessing
Notify	29	mstevens	AST_L1BT	001	Dec 31 2001 11:42AM	Dec 31 2001 11:42AM	FinishedActionProcessing
Acquire	29	mstevens	AST_L1BT	001	Dec 31 2001 11:42AM	Dec 31 2001 11:42AM	FinishedActionProcessing
Notify	35	kcarr	MOD021KM	001	Jan 2 2002 3:12PM	Jan 2 2002 3:12PM	FinishedActionProcessing
Notify	35	kcarr	MOD021KM	001	Jan 2 2002 1:57PM	Jan 2 2002 1:57PM	FinishedActionProcessing
Notify	35	kcarr	MOD021KM	001	Jan 2 2002 1:49PM	Jan 2 2002 1:49PM	FinishedActionProcessing

Figure 56. NSBRV List Action Queue page

The remaining major page is the **List Statistics** page, shown in Figure 57. It provides summary information concerning the processing of events and actions related to subscriptions. Using this page, the operator can monitor subscription processing activity, such as numbers of notifications and actions, total and average times for notifications and acquires, and other information. This page reflects processing based on what is currently in the database.

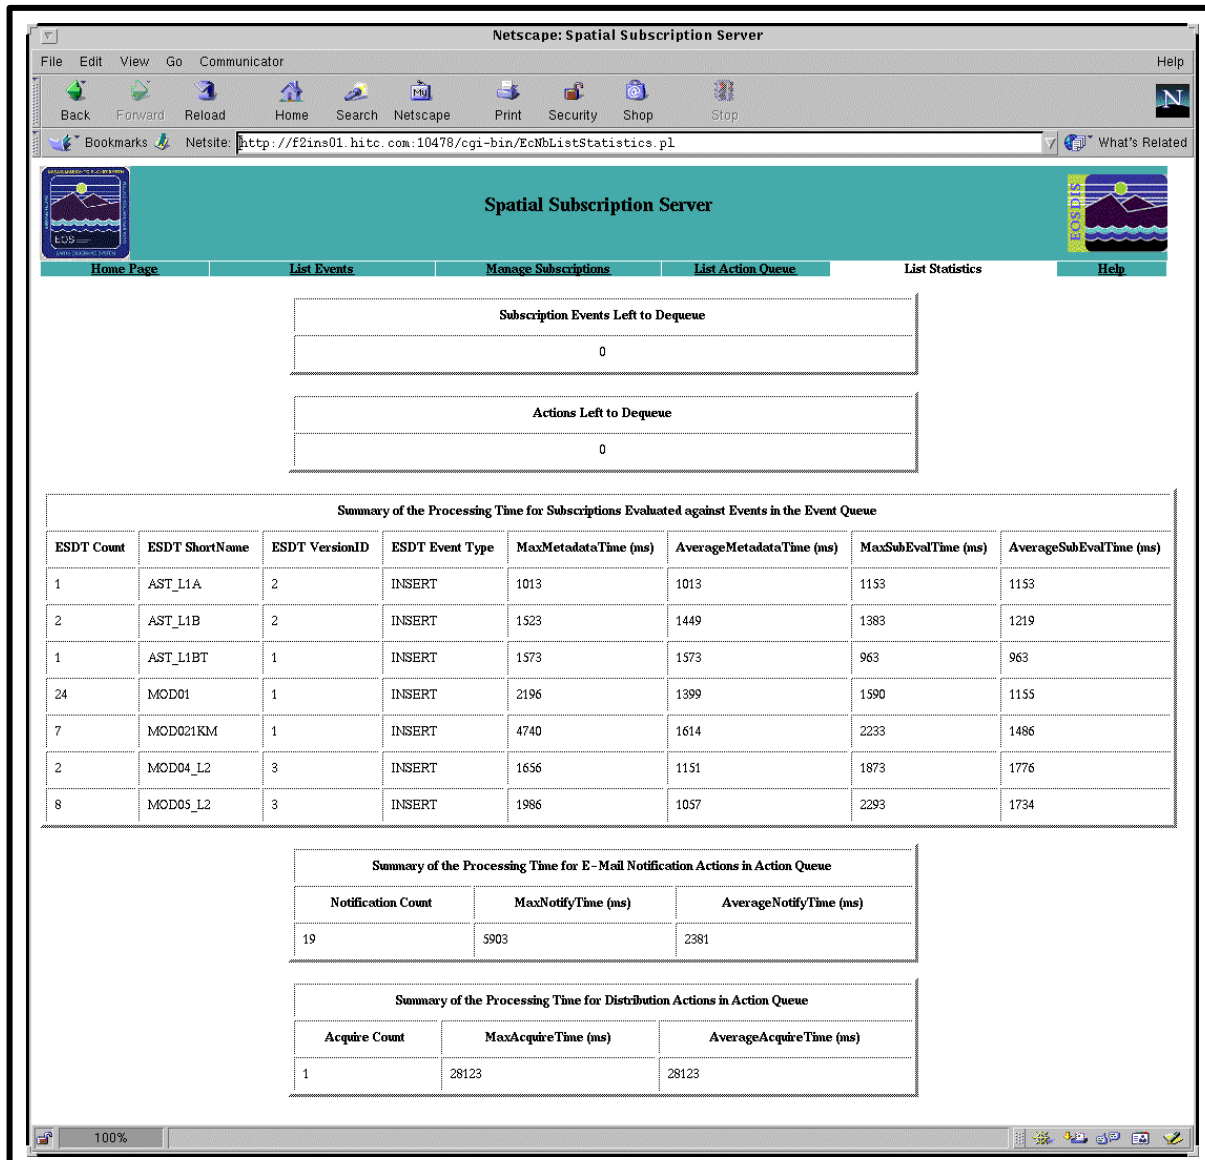


Figure 57. NSBRV List Statistics page

Working with Existing Subscriptions

This section examines how you can use the NSBRV GUI for certain tasks in managing subscriptions. The first thing you will need to do for many functions is to launch the GUI. The procedure for launching the GUI is provided separately here and referenced in other procedures.

Launch the NSBRV GUI

- 1 At the UNIX command shell prompt, type **setenv DISPLAY *clientname*:0.0** and then press the **Return/Enter** key.
 - For *clientname*, use either the local terminal/workstation IP address or its machine name.
 - 2 Start the log-in to a Netscape host by typing **/tools/bin/ssh *hostname*** (*e.g.*, g0ins02, e0ins02, l0ins02, n0ins02) at the UNIX command shell prompt, and press the **Return/Enter** key.
 - If you receive the message, **Host key not found from the list of known hosts. Are you sure you want to continue connecting (yes/no)?** type yes (“y” alone does not work).
 - If you have previously set up a secure shell passphrase and executed **sshremote**, a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears; continue with Step 3.
 - If you have not previously set up a secure shell passphrase; go to Step 4.
 - 3 If a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears, type your *Passphrase* and then press the **Return/Enter** key. Go to Step 5.
 - 4 At the **<user@remotehost>'s password:** prompt, type your *Password* and then press the **Return/Enter** key.
 - You are logged in and a UNIX command shell prompt is displayed.
 - 5 Type **netscape** and then press the **Return/Enter** key.
 - The Netscape web browser is displayed.
 - 6 Click in the **Netsite:** field.
 - The field is highlighted.
 - 7 Type the Universal Resource Locator (URL) for the NSBRV GUI and then press the **Return/Enter** key.
 - The NSBRV **Home Page** is displayed, offering access to NSBRV functions (**List Events**, **Manage Subscriptions**, **List Action Queue**, and **List Statistics**) as well as a tab for **Help** in navigating the GUI.
-

Suppose that you now want to view subscriptions related to MOD01 data, and then to review detailed information about a particular subscription. To obtain the list you will need to use the NSBRV **Manage Subscriptions** tab, filter on **Collection**, and then select and apply an option to **View** the particular subscription. Use the following procedure.

Use the NSBRV GUI to List and View Subscriptions and Themes in the NSBRV Database

- 1** Launch the NSBRV GUI (refer to procedure **Launch the NSBRV GUI**, previous section of this lesson).
- 2** Click on the **Manage Subscriptions** tab.
 - The Manage Subscriptions page is displayed with a table listing all subscriptions in the NSBRV database.
 - The column headers in the table, except for **Version** and **Event Type**, are links for sorting the list. There are also buttons for filtering the list.
- 3** To filter the list for display of only MOD01 subscriptions, click on the **Collection** option button and then scroll and click to select **MOD01**.
 - The selected choice is displayed on the option button and a list is displayed showing only MOD01 subscriptions.
- 4** It may be desirable to filter the list. For example, if you are looking for a subscription created for a specific user, click on the **User** option button and then scroll and click to select the appropriate User Id.
 - The selected choice is displayed on the option button and a list is displayed showing only subscriptions for the selected User Id.
- 5** Scroll if necessary to reach the sought subscription and click on the **View** radio button for that subscription in the last column.
 - The button is filled to indicate selection of the **View** option.
- 6** Click on the **Apply** button for the subscription.
 - Detailed information is displayed for the selected subscription.
- 7** From the page displaying detailed information for a selected subscription, or from the Manage Subscriptions page obtained in step 2, click on the **List Themes** tab.
 - The GUI displays a prompt **Enter first few characters of theme name (or leave blank to view all):** and provides an entry field and an **Apply** button.
- 8** If you know the name of the theme you are seeking, click in the entry field and type the first three or four characters of the name; otherwise go to step 9.
 - Any typed entry appears in the field.

9 Click on the **Apply** button.

- The GUI displays a list of themes, either those whose names start with the letters typed in step 8, or all those in the database, with an **Apply** button for each theme for listing subscriptions associated with that theme.

10 To list the subscriptions associated with a given theme, click the **Apply** button for that theme.

- The GUI lists the subscriptions associated with the selected theme, providing the same information and options for each subscription as provided in the results of step 2 (**View**, **Update**, and **Cancel**). It also provides the option to **Suspend All**, **Resume All**, or **Cancel All** the subscriptions for that theme.

11 To view one of the subscriptions, repeat steps 5 and 6.

- Detailed information is displayed for the selected subscription.
-

It may become desirable to extend the period of retention in a Data Pool insert subscription. For example, if conditions arise that increase user interest in the data that will be inserted as a result of the subscription, keeping the data in the Data Pool longer is a way to be responsive to that user interest. Such conditions might be unusual weather activity, flooding, volcanic activity, or other events in the area covered by the subscription. Suppose that Dr. L. Abuser contacts the DAAC with a request that, for the remainder of the two-year period of the Data Pool insert subscription placed in response to the original request, the data placed in the data pool be kept for 60 days instead of 30. Use the following procedure to extend the retention period.

Use the NSBRV GUI to Extend the Period of Retention in a Data Pool Insert Subscription

1 Launch the NSBRV GUI (refer to procedure **Launch the NSBRV GUI**, previous section of this lesson).

- **Note:** At various points in this procedure, you may encounter a security information warning. Unless you know of a specific potential danger that you must avoid, click on the **Continue submission** button when this warning is displayed.

2 Click on the **Manage Subscriptions** tab.

- The **Manage Subscriptions** page is displayed with a table listing all subscriptions in the NSBRV database.

3 Click on the **User** option button and then scroll and click to select the appropriate User Id.

- The selected choice is displayed on the option button and a list is displayed showing only subscriptions for the selected User Id.

- 4 Scroll if necessary to reach the sought subscription and click on the **Update** radio button for that subscription in the last column.
 - The button is filled to indicate selection of the **Update** option.
- 5 Click on the **Apply** button for the subscription.
 - A page for the subscription is displayed with a block of option buttons and fields to permit updating qualifiers, **Latitude** and **Longitude** coordinates, and **Actions** information. It may be necessary to use a scroll bar to reveal all entry areas.
 - At the bottom, there is an **Update Subscription** button for submitting the changes.
- 6 To change the retention period, click at the end of the **Retention Period** field in the **Data Pool Information** portion of the **Actions** block and backspace to delete the displayed number (for this exercise, **30**).
 - The cursor is displayed in the field.
- 7 Type the desired extended retention period (within any Data Pool constraints imposed by the DAAC) (in this case, type **60**).
 - The typed entry is displayed in the field.
- 8 Click on the **Update Subscription** button.
 - A confirmation screen is displayed, a message **Are you sure that you are ready to update?** appears above an **Update Subscription** button, and a link permits the operator to **Return to Previous Page**.
- 9 Click on the **Update Subscription** button.
 - A message is displayed confirming that the subscription was updated.

Using the NSBRV GUI to Manage Bundling Orders

The NSBRV GUI is used to list, view, create, update, list associated subscriptions, and cancel bundling orders. These functions are accessible through the **Manage Bundling Orders** tab. The following procedures are applicable.

Use the NSBRV GUI to Obtain a List of Bundling Orders and View a Bundling Order

- 1 Launch the NSBRV GUI (refer to procedure **Launch the NSBRV GUI**, previous section of this lesson).
 - **Note:** At various points in this procedure, you may encounter a security information warning. Unless you know of a specific potential danger that you must avoid, click on the **Continue submission** button when this warning is displayed.

- 2 Click on the **Manage Bundling Orders** tab.
 - The **Manage Bundling Orders** page is displayed with a table listing all bundling orders in the NSBRV database, providing for each a **Bundling Order ID**, the **User ID** of the person who created the bundling order, the **Creation Date**, **Expiration Date**, **Media Type**, and **Status** of the bundling order. The column headers for these data are links for sorting the list on the information in the column.
 - The page also provides for each bundling order the options to **View**, **Update**, **Cancel**, and **List Subs**, with an **Apply** button to implement the selected option.
 - 3 To list the subscriptions associated with a bundling order, click on the radio button next to **List Subs** in the **Choose Bundling Order Action** at the end of the row for the bundling order.
 - The radio button is filled to indicate selection.
 - 4 To implement the selected action, click on the **Apply** button at the end of the row.
 - The GUI displays a page listing the subscriptions associated with the selected bundling order, providing for each associated subscription the information identified in step 2.
 - 5 Scroll if necessary to reach the sought bundling order and click on the **View** radio button for that bundling order in the last column.
 - The button is filled to indicate selection of the **View** option.
 - 6 Click on the **Apply** button at the end of the row for the bundling order.
 - Detailed information is displayed for the selected bundling order.
-

To cancel a bundling order, use the following procedure.

Use the NSBRV GUI to Cancel a Bundling Order and Its Associated Subscriptions

- 1 Launch the NSBRV GUI (refer to procedure **Launch the NSBRV GUI**, previous section of this lesson).
 - **Note:** At various points in this procedure, you may encounter a security information warning. Unless you know of a specific potential danger that you must avoid, click on the **Continue submission** button when this warning is displayed.

- 2 Click on the **Manage Bundling Orders** tab.
 - The **Manage Bundling Orders** page is displayed with a table listing all bundling orders in the NSBRV database, providing for each a **Bundling Order ID**, the **User ID** of the person who created the bundling order, the **Creation Date**, **Expiration Date**, **Media Type**, and **Status** of the bundling order. The column headers for these data are links for sorting the list on the information in the column.
 - The page also provides for each bundling order the options to **View**, **Update**, **Cancel**, and **List Subs**, with an **Apply** button to implement the selected option.
 - 3 To cancel a bundling order, click on the radio button next to **Cancel** in the **Choose Bundling Order Action** at the end of the row for the bundling order to be canceled.
 - The radio button is filled to indicate selection.
 - 4 To implement the selected action, click on the **Apply** button at the end of the row.
 - A confirmation message asks **Are you sure that you want to cancel the bundling order? Any associated subscriptions will be simultaneously canceled..**
 - 5 Click on the **Yes** button.
 - The GUI displays a confirmation that **Bundling Order <nnnnnn> was canceled.**
-

To add a bundling order, use the following procedure.

Use the NSBRV GUI to Add a Bundling Order

- 1 Launch the NSBRV GUI (refer to procedure **Launch the NSBRV GUI**, previous section of this lesson).
 - **Note:** At various points in this procedure, you may encounter a security information warning. Unless you know of a specific potential danger that you must avoid, click on the **Continue submission** button when this warning is displayed.
- 2 Click on the **Manage Bundling Orders** tab.
 - The **Manage Bundling Orders** page is displayed with a table listing all bundling orders in the NSBRV database, providing for each a **Bundling Order ID**, the **User ID** of the person who created the bundling order, the **Creation Date**, **Expiration Date**, **Media Type**, and **Status** of the bundling order. The column headers for these data are links for sorting the list on the information in the column.
 - The page also provides for each bundling order the options to **View**, **Update**, **Cancel**, and **List Subs**, with an **Apply** button to implement the selected option.
- 3 Click on the **Add Bundling Order** link.
 - The GUI displays the **Add Bundling Order** page.

- 4 Click in the **User ID** field to place the cursor in the field.
 - The cursor is displayed in the **User ID** field.
- 5 Type your **User ID**.
 - The typed entry is displayed in the **User ID** field.
- 6 If it is necessary or desirable to change the displayed default **Expiration Date**, click in the **Expiration Date** field and change to the desired date, either by dragging the cursor to highlight the current value and typing over it or by backspacing to delete the current value and typing the desired date.
 - The desired date is displayed in the **Expiration Date** field.
- 7 Click on the **Media Type** option button to display the **Media Type** options and click on the desired option to select it.
 - The desired **Media Type** option is displayed on the option button.
- 8 Click on the **continue** button.
 - A confirmation dialog box asks **Your present values have been entered. Continue?.**
- 9 Click on the **OK** button in the confirmation dialog box.
 - The GUI displays a data page appropriate for the selected **Media Type**.
 - For hard media, the page provides fields to specify shipping information.
 - For FTP Push, the page provides fields to specify FTP Push parameters.
 - For FTP Pull, the page does not require media parameters.
 - The displayed data page contains fields to specify an **Email Address** and any additional **User String**. It also provides an option button to select a **Distribution Priority** (**VHIGH**, **HIGH**, **NORMAL**, **LOW**, or **XPRESS**).
 - The displayed data page provides fields to specify completion criteria for a bundle (with defaults appropriate to the selected **Media Type**).
 - Minimum Bundle Size (GB).
 - Minimum Granule Count.
 - Maximum Bundle Age (days).

- 10 For each required field (identified on the GUI by an asterisk) and each field where a change is desirable, click in the field and change to the desired entry. This may entail typing a value in a blank field or changing a default, either by dragging the cursor to highlight the current value and typing over it or by backspacing to delete the current value and typing the desired value.
 - The fields display entered and/or desired values.
 - 11 Click on the **Add Bundling Order** button.
 - The GUI displays a confirmation that **Bundling Order <nnnnnn> was created.**
-

Use the following procedure to update a bundling order.

Use the NSBRV GUI to Update a Bundling Order

- 1 Launch the NSBRV GUI (refer to procedure **Launch the NSBRV GUI**, previous section of this lesson).
 - *Note:* At various points in this procedure, you may encounter a security information warning. Unless you know of a specific potential danger that you must avoid, click on the **Continue submission** button when this warning is displayed.
- 2 Click on the **Manage Bundling Orders** tab.
 - The **Manage Bundling Orders** page is displayed with a table listing all bundling orders in the NSBRV database, providing for each a **Bundling Order ID**, the **User ID** of the person who created the bundling order, the **Creation Date**, **Expiration Date**, **Media Type**, and **Status** of the bundling order. The column headers for these data are links for sorting the list on the information in the column.
 - The page also provides for each bundling order the options to **View**, **Update**, **Cancel**, and **List Subs**, with an **Apply** button to implement the selected option.

- 3 To update a bundling order, click on the radio button next to **Update** in the **Choose Bundling Order Action** at the end of the row for the bundling order.
 - The radio button is filled to indicate selection.
- 4 To implement the selected action, click on the **Apply** button at the end of the row.
 - The GUI displays the **Update Bundling Order <nnnnnn>** page, showing the **User Id**, **Expiration Date**, and **Media Type**.
- 5 If it is necessary or desirable to update the **Expiration Date**, click in the **Expiration Date** field and change to the desired date, either by dragging the cursor to highlight the current value and typing over it or by backspacing to delete the current value and typing the desired date.
 - The desired date is displayed in the **Expiration Date** field.
- 6 If it is necessary or desirable to change the media type, click on the **Media Type** option button to display the **Media Type** options and click on the desired option to select it.
 - The desired **Media Type** option is displayed on the option button.
- 7 Click on the **continue** button.
 - A confirmation dialog box asks **Your present values have been entered. Continue?**
- 8 Click on the **OK** button in the confirmation dialog box.
 - The GUI displays a data page appropriate for the selected **Media Type**.
 - For hard media, the page provides fields to specify shipping information.
 - For FTP Push, the page provides fields to specify FTP Push parameters.
 - For FTP Pull, the page does not require media parameters.
 - The displayed data page contains fields to specify an **Email Address** and any additional **User String**. It also provides an option button to select a **Distribution Priority** (**VHIGH**, **HIGH**, **NORMAL**, **LOW**, or **XPRESS**).
 - The displayed data page provides fields to specify completion criteria for a bundle (with defaults appropriate to the selected **Media Type**).
 - Minimum Bundle Size (GB).
 - Minimum Granule Count.
 - Maximum Bundle Age (days).
- 9 For each required field (identified on the GUI by an asterisk) and each field where a change is desirable, click in the field and change to the desired entry. This may entail typing a value in a blank field or changing a default, either by dragging the cursor to highlight the current

value and typing over it or by backspacing to delete the current value and typing the desired value.

- The fields display entered and/or desired values.

10 Click on the **Update Bundling Order** button.

- The GUI displays a confirmation that **Bundling Order <nnnnnn> was successfully updated.**
-

Using the NSBRV GUI for Monitoring the Spatial Subscription Server

The NSBRV GUI provides pages that can be used to keep track of current actions as the NSBRV processes Acquires and Notifications, and to track NSBRV performance. These capabilities are accessible through tabs on the NSBRV **Home Page**.

The **List Action Queue** tab is used to view Acquire and Notification actions being processed by the NSBRV. The following procedure is applicable.

Use the NSBRV GUI to View the Acquire and Notification Actions Being Processed

- 1** Launch the NSBRV GUI (refer to procedure **Launch the NSBRV GUI**, previous section of this lesson).
 - 2** Click on the **List Action Queue** tab.
 - The **List Action Queue** information is displayed in a table listing acquire and notification actions that are being processed. On this page, the **Action Type** and **Subscription Id** column headers are links for sorting the list, and there are also **Action Type**, **Subscription**, and **Status** option buttons and a filter button for filtering the list.
-

The **List Statistics** tab is used to review summary information concerning the processing of events and subscriptions. The following procedure is applicable.

Use the NSBRV GUI to View Statistics on NSBRV Processing of Events and Actions

- 1** Launch the NSBRV GUI (refer to procedure **Launch the NSBRV GUI**, previous section of this lesson).
- 2** Click on the **List Statistics** tab.

- The **List Statistics** information is displayed, showing **Subscription Events Left to Dequeue**, **Actions Left to Dequeue**, a **Summary of Subscription Events in SubEventQueueLog**, a **Summary of Notification Actions in ActionQueueLog**, a **Summary of Acquire Actions in ActionQueueLog**, **Details of Subscription Events in SubEventQueueLog**, **Details of Notification Actions in ActionQueueLog**, and **Details of Acquire Actions in ActionQueueLog**.
-

Working with Data Pool Scripts

There are several Data Pool scripts that provide the operator with utilities or applications for managing Data Pool maintenance. These include:

- **Update Granule Utility:** a script to update granule expiration (extend the period of retention) and, optionally, retention priority, for selected science granules already in the Data Pool. For Synergy III, the utility allows operators to extend the expiration of all granules associated with a particular thematic collection or upgrade their expiration priority.
- **Data Pool Cleanup Utility:** a script to remove expired granules from Data Pool disks and inventory and free up space in the Data Pool. For Synergy III, it allows operators to clean up a thematic collection. It permits unlinking granules from a thematic collection, and permits data pool removal of granules that belong to a given thematic collection if they meet the cleanup criteria specified by the other command line parameters.
- **Data Pool Access Statistics Utility (DPASU):** scripts for processing available Data Pool access logs to extract and summarize statistics on FTP and web access to data in the Data Pool. The statistics are stored in the Data Pool database to be used for producing tabular reports that can be loaded into a spreadsheet program for sorting, graphing, or other manipulation.
- **Data Pool Archive Access Statistics Data Utility:** an operational support tool for archiving the granule access data from the Data Pool database into a tab-delimited ASCII file.
- **Data Pool Delete Access Statistics Data Utility:** an operational support tool for deleting granule access data in a specified time range from the Data Pool database.
- **Data Pool Restore Access Statistics Data Utility:** an operational support tool for restoring granule access data for a specific time range from an ASCII archive file to the Data Pool database.
- **Batch Insert Utility:** a command-line utility for inserting non-ECS data and ECS data that are already in the archive into the Data Pool.

Extending the Period of Retention for Granules in the Data Pool

We have seen that a change in user interest in data from a particular location may arise because of unusual circumstances (e.g., weather, natural event) and that as a result it may be desirable to extend the period of retention in a Data Pool insert subscription. Such circumstances may also make it desirable to retain certain data already in the Data Pool for a longer period of time than originally specified. Data Pool maintenance personnel can run the Update Granule Utility to update the expiration date for selected science granules. This utility also permits modifying a granule's retention priority, which can affect how soon the Data Pool Cleanup Utility removes the granule from the Data Pool.

The Update Granule Utility permits updating granule information using a command-line interface. The following options may be used:

- noprompt**: suppressing prompts and detailed information display.
- theme**: specifies a valid theme name (i.e., a character string that matches an existing theme name in the Data Pool inventory).

A single granule may be updated using manual input. Multiple granule updates can be handled using an input file containing a list of granules to be updated, or by specifying a theme. The input file must be structured as a list of granules to be processed, one per line. Each line contains a granule ID (reflecting the Sybase entry in the Data Pool database), an expiration date, and (optionally) a new retention priority, the value of which may be null (i.e., left blank). The fields are separated by a single space. There should be no blank lines before the first or after the last granule in the list. The file contents should be similar to the following example.

```
GRANULE_ID_4832 EXP_DATE=2002/2/28 RETENTION=255
GRANULE_ID_4876 EXP_DATE=2002/2/28 RETENTION=200
GRANULE_ID_4883 EXP_DATE=2002/2/28 RETENTION=
GRANULE_ID_4937 EXP_DATE=2002/2/28
GRANULE_ID_4966 EXP_DATE=2002/2/28 RETENTION=255
```

When updating the granules associated with a theme, the utility updates the expiration date of a granule associated with that theme if and only if the new expiration date specified is later than the current expiration date of the granule. It updates the retention priority of a granule associated with that theme if and only if the new expiration priority specified is higher than the current retention priority of the granule.

The Update Granule Utility connects to the Data Pool database and calls Sybase stored procedures to perform the requested updates. Therefore, the utility runs only if the Data Pool database server is running and if the database is available. It also assumes the stored procedures are present. The Granule Update Utility may be run as a background process, with suppression of all warning/error messages and confirmation prompts if desired. When the utility is run, it writes information, any warnings, any errors, and messages to a log file about granules as they are updated.

Assume that a user contacts the DAAC with a request to update (extend) the expiration date to the end of February 2002 for selected granules in the Data Pool, and provides a list of granule IDs for the selected granules. The following procedure is applicable.

Use the Update Granule Utility to Extend the Retention for Selected Science Granules

- 1 Log in at the machine on which the Update Granule Utility is installed (e.g., e0dps01, g0dps01, l0dps01, n0dps01).
- 2 To change to the directory for starting the Update Granule Utility, type **cd /usr/ecs/<MODE>/CUSTOM/utilities** and then press the **Return/Enter** key.
 - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/utilities**.
- 3 At the UNIX prompt, enter the command to start the Update Granule Utility, in the form **EcDIUpdateGranule.pl <command line parameters>**. For this exercise, use the following command:

EcDIUpdateGranule.pl <MODE> -file tr_list

(*Note:* The first command-line parameter specified must be **<MODE>**, a valid, existing Data Pool mode [e.g., OPS, TS1, TS2]).

- *Note:* The following six permutations are valid command-line entries for initiating the Update Granule utility:
 - **EcDIUpdateGranule.pl <MODE> -file <filename>** (to update granules listed in an input file named **<filename>** while displaying all summary information to the operator, and asking confirmation of the update).
 - **Ec DIUpdateGranule.pl <MODE> -grnid <granuleID> -exp <expiration date> [-ret <retention priority>]** (to update a granule identified by its **<granuleID>** with a new expiration date and, optionally, a new retention priority while displaying all summary information to the operator, and asking confirmation of the update).
 - **EcDIUpdateGranule.pl <MODE> -noprompt -file <filename>** (to update granules listed in an input file named **<filename>** with no confirmation or information displayed to the operator).
 - **EcDIUpdateGranule.pl <MODE> -noprompt -grnid <granuleID> -exp <expiration date> [-ret <retention priority>]** (to update a granule identified by its **<granuleID>** with a new expiration date and, optionally, a new retention priority with no confirmation or information displayed to the operator).
 - **EcDIUpdateGranule.pl <MODE> -theme <themename> -exp <expiration date> [-ret <retention priority>]** (to update a granule identified by its **<themename>** with a new expiration date and, optionally, a new retention priority while displaying all summary information to the operator, and asking confirmation of the update).
 - **EcDIUpdateGranule.pl <MODE> -noprompt -theme <themename> -exp <expiration date> [-ret <retention priority>]** (to update a granule identified by

its *<themename>* with a new expiration date and, optionally, a new retention priority with no confirmation or information displayed to the operator.

- The utility executes and displays a confirmation prompt similar to the following:

```
You are about to start updating granules.
-----
Total number of granules: 11
Total size of granules: 8.61339673772454 MB
Do you wish to continue processing the update? [y/n]y
```

4 Type **y** and then press the **Return/Enter** key.

- The utility completes execution and displays output similar to the following:

```
Update completed.
Please check the database to ensure proper completion.

Update took 2 seconds to complete

Gracefully exiting...
```

- To check the database, have the Database Administrator use *isql* commands on the Data Pool database host to query the DIGranuleExpirationPriority table. It may also be useful to examine the Update Granule Utility log file to determine whether there were any problems with the execution. To examine that log file, go to Steps 5 and 6.

5 To change to the directory containing the Update Granule Utility log file and other log files, type **cd /usr/ecs/<MODE>/CUSTOM/logs**, and then press the **Return/Enter** key.

- The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/logs**.

6 To examine the Update Granule Utility log file, type **pg EcDIUpdateGranule.log** and then press the **Return/Enter** key.

- The first page of the log file is displayed; additional sequential pages can be displayed by pressing the **Return/Enter** key at the **:** prompt. It is also possible to search forward by typing **/<search item>**. For example, to search the log file for reference to one of the granules updated, type **/<granuleID>** and then press the **Return/Enter** key.
- Although this procedure is written for the **pg** command, any UNIX editor or visualizing command (e.g., **vi**, **view**, **more**, **tail**) can be used to review the log.
- The log entries have a time and date stamp; about the time that the update was executed, the log should show entries similar to the following:

```
2001/11/29 15:52:50.814:Update started...
```

```
2001/11/29 15:52:50.964:Granule 4871 updated
2001/11/29 15:52:51.083:Granule 4954 updated
2001/11/29 15:52:51.212:Granule 4955 updated
2001/11/29 15:52:51.346:Granule 4956 updated
2001/11/29 15:52:51.409:Granule 4957 updated
2001/11/29 15:52:51.688:Granule 4959 updated
2001/11/29 15:52:51.778:Granule 4961 updated
2001/11/29 15:52:51.998:Granule 4963 updated
2001/11/29 15:52:52.107:Granule 4963 updated
2001/11/29 15:52:52.394:Granule 4964 updated
2001/11/29 15:52:52.569:Granule 4966 updated
2001/11/29 15:52:52.590:Update ended.
```

```
2001/11/29 15:52:52.608:This update took approximately 2 seconds
```

- If the log indicates errors or warnings, it may be necessary to correct the condition identified in the entry (e.g., edit the data in the granule list in the input file) and run the utility again. Specific error entries depend on the error that occurred; examples of error entries in the log may be similar to the following:

```
4959      AST_04      1 0.03962299 Jul 30 2001 12:00AM Feb  2 1998
11:59PM      255      2
```

```
Warning: The new expiration date for the above granule is less than
or equal to today's date.
```

```
DATABASE ERROR:Server message number=120001 severity=16 state=1
line=33      server=f2acg01_srvr      procedure=ProcSelectGrExpiration
text=ProcSelectGrExpiration: Requested granule id not in database.
```

```
2001/11/29 15:50:36.647:Sybase Lookup ==> ERRORS WERE FOUND WITH
GRANULE "4654". (It may not exist or contains the wrong format).
```

```
2001/11/29 15:50:36.663:
```

```
EcDlUpdateGranule_1.pl aborted due to insufficient processing data:
All the granule triplets had errors.
```

Running the Data Pool Cleanup Utility

The Data Pool Cleanup Utility permits ECS Operations Staff to remove expired granules from the Data Pool disks and corresponding inventory. It must be executed on the machine where the granules are located. Qualification for cleanup is based on three criteria: expiration date/time, retention priority, and theme association. Figure 58 illustrates how the utility uses expiration date/time and retention priority in combination to select granules for removal from the Data Pool. The operator specifies values for these criteria and, if desired, themes and cross references, to control cleanup operations. It is also possible to use the utility to validate the Data Pool inventory and disk cache by checking for and removing *orphans* (disk files not associated with any entry in the Data Pool inventory) and *phantoms* (entries in the Data Pool inventory that refer to files that no longer exist on the Data Pool disks).

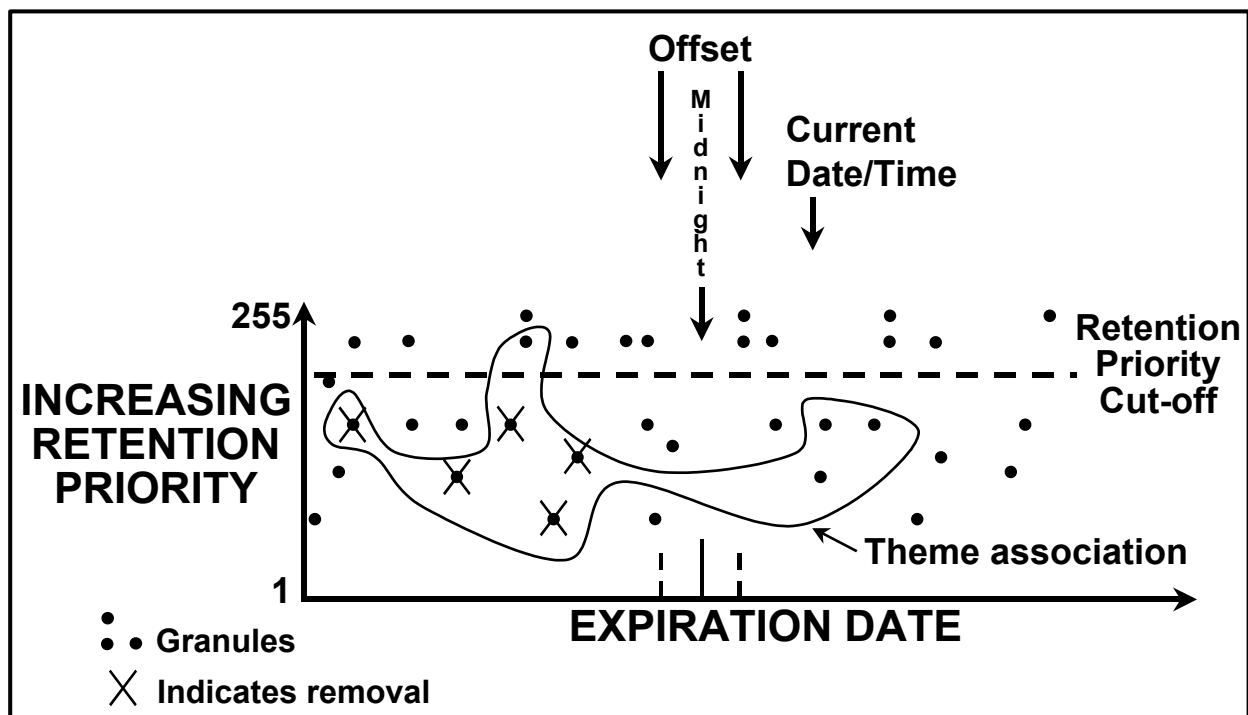


Figure 58. The Data Pool Cleanup Utility uses Expiration Date, Retention Priority, and/or Theme Association to qualify granules, marked with X, for deletion

To determine whether a granule qualifies for deletion, the utility first compares the granule's expiration date (insert date plus retention period in days specified in the insert subscription) with a cut-off date/time. If a granule's expiration date is prior to the cut-off, the granule qualifies as expired.

- The default cut-off date/time is set to midnight of the previous day.
- The operator is permitted to specify an 'offset' in hours to add or subtract hours from the previous midnight to determine a cut-off date/time for deletion.

Next, the utility compares the granule's retention priority with any priority limit the operator has specified to identify those granules that should be retained in the Data Pool even though their expiration date has passed.

- Retention priority is an integer from 1 to 255.
- Retention priority for granules already in the Data Pool may be modified using the granule expiration update script.

If the operator specifies a theme name, the utility applies the removal criteria only to those granules associated with that theme.

As suggested by Figure 58, the Data Pool Cleanup Utility removes those granules with expiration date prior to the cut-off date/time and with retention priority is less than or equal to the specified limit. If a priority limit is not specified in command-line input parameters at the time it is invoked, the Cleanup Utility reads the parameter 'DEFAULT_LIMIT' from its

configuration file to get a priority limit. If the operator does not wish to use retention priority as a criterion for deletion, the default limit should be set to 255.

The Cleanup Utility can alternatively take as input a file listing granuleId's for granules to be deleted. The file can contain single or multiple granuleId's per line separated by whitespace.

If the Cleanup Utility is interrupted during execution, upon restart it continues from the point of interruption. Further, in the interest of low database contention, the Cleanup Utility allows only one instance of itself to execute. The Cleanup Utility also provides an option to suppress operator prompts/messages, although there are not many. Upon completion of removing the granule files from the Data Pool disks, the Cleanup Utility determines if there is sufficient free space to update the 'NoFreeSpace' flag in the Data Pool database, if necessary.

Cleanup Associated with Themes

The Cleanup Utility cleans up non-ECS data just as it does ECS data. It can remove granule cross references associated with a given theme, and also remove the granules associated with the theme. The option **-themexref** specifies a theme for which all cross references are to be removed from the Data Pool. The option **-theme** specifies a theme for which associated granules are to be removed. If a granule is referenced to more than one theme, the **-theme** option removes only the cross reference to the specified theme, without removing the granule.

Cleanup of Data Pool/Inventory Discrepancies ("Validation")

The Cleanup Utility provides a validation capability to identify and clean up discrepancies between the Data Pool inventory and disk content. It may also be used just to log such discrepancies. This capability uses the options **-orphan** (to find/remove data in the Data Pool that is not represented by entries in the Data Pool inventory) and **-phantom** (to find/remove entries in the Data Pool inventory that have one or more science or metadata files, or associated browse files, missing from the Data Pool. To specify just logging of the discrepancies, the operator uses the option **-nofix**.

Actions Performed by the Cleanup Utility

As part of its processing, the Cleanup Utility performs the following actions:

- removes from the Data Pool and inventory all data pool granules, browse files, and browse links that meet the specified cleanup criteria (provided that no other granules are cross-referenced to them – i.e., linked by a theme).
- exports a list of deleted granules for accessibility by the EOS Clearing House (ECHO); to do this, the utility invokes an external utility, EcOsBulkURL, which generates an XML file containing a list of deleted granules that qualify for ECHO export and stores it in the directory /datapool/<MODE>/user/URLExport.

- removes all HEG conversion files associated with the HEG order IDs that have the status of “DONE” or “FAILED” and a timestamp older than the cleanup age specified by the **HEGCleanupAge** parameter in the DIConfig table of the Data Pool database. (HEG orders and conversion files are generated when end users request HEG-converted data using the Data Pool Web Access tool.)
- (if specified) checks the Data Pool inventory and disk content for the existence of orphans and/or phantoms, removing or just logging them depending on the command-line options specified.
- determines how much free space has been cleared from the Data Pool disks and, if that space is greater than or equal to the amount of disk space specified in the MIN_FREE_SPACE parameter, updates the ‘NoFreeSpaceFlag’ in the Data Pool database, setting it to permit insertion of additional data.

The utility may be executed using a **-noprompt** argument to suppress all confirmations and warnings normally displayed to standard output. To initiate the Data Pool Cleanup Utility manually, use the following procedure.

Invoke the Data Pool Cleanup Utility Manually

- 1 Log in at the machine where the Data Pool Cleanup Utility is installed (e.g., e0dps01, g0dps01, l0dps01, n0dps01).
 - **Note:** The operator who is executing the script must have the privilege to remove science, metadata, and browse files from the Data Pool disks.
- 2 To change to the directory for starting the Data Pool Cleanup Utility, type **cd /usr/ecs/<MODE>/CUSTOM/utilities** and then press the **Return/Enter** key.
 - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/utilities**.
- 3 At the UNIX prompt, enter the command to start the Cleanup Utility, in the form **EcDlCleanupDataPool.pl <MODE> [-option1 <value1> ... -optionN <valueN>]**, specifying up to three options. For this exercise, use the following command:

EcDlCleanupDataPool.pl <MODE>

- **Note:** The following are examples of valid command-line entries for initiating the Data Pool Cleanup utility:
 - **EcDlCleanupDataPool.pl <MODE>** (to delete all granules with retention priority less than or equal to the configured default limit that have expiration dates before midnight of the previous day).
 - **EcDlCleanupDataPool.pl <MODE> -noprompt** (to delete granules with retention priority less than or equal to the configured default limit that have expiration dates before midnight of the previous day while suppressing all operator prompts and confirmations).

- **EcDlCleanupDataPool.pl** *<MODE>* **-limit** *<priority limit>* (to delete all granules with retention priority less than or equal to the specified *<priority limit>* that have expiration dates before midnight of the previous day).
- **EcDlCleanupDataPool.pl** *<MODE>* **-limit** *<priority limit>* **-noprompt** (to delete all granules with retention priority less than or equal to the specified *<priority limit>* that have expiration dates before midnight of the previous day while suppressing all operator prompts and confirmations).
- **EcDlCleanupDataPool.pl** *<MODE>* **-offset** *<±hours>* (to delete granules with retention priority less than or equal to the default limit that have expiration dates before midnight of the previous day plus or minus the specified number of *<hours>* -- e.g., **EcDlCleanupDataPool.pl OPS -offset -5** deletes granules with expiration date before 7:00 p.m. yesterday).
- **EcDlCleanupDataPool.pl** *<MODE>* **-offset** *<±hours>* **-noprompt** (to delete granules with retention priority less than or equal to the default limit that have expiration dates before midnight of the previous day plus or minus the specified number of *<hours>* while suppressing all operator prompts and confirmations).
- **EcDlCleanupDataPool.pl** *<MODE>* **-offset** *<±hours>* **-limit** *<priority limit>* (to delete granules with retention priority less than or equal to the specified *<priority limit>* that have expiration dates before midnight of the previous day plus or minus the specified number of *<hours>*).
- **EcDlCleanupDataPool.pl** *<MODE>* **-offset** *<±hours>* **-limit** *<priority limit>* **-noprompt** (to delete granules with retention priority less than or equal to the specified *<priority limit>* that have expiration dates before midnight of the previous day plus or minus the specified number of *<hours>* while suppressing all operator prompts and confirmations).
- **EcDlCleanupDataPool.pl** *<MODE>* **-file** *<filename>* (to delete granules listed by granuleID in the input file named *<filename>*, and any associated browse granules). *Note:* The **-file** option may not be used with the **-offset** option or **-limit** option.
- **EcDlCleanupDataPool.pl** *<MODE>* **-file** *<filename>* **-noprompt** (to delete granules listed by granuleID in the input file named *<filename>*, and any associated browse granules while suppressing all operator prompts and confirmations).
- **EcDlCleanupDataPool.pl** *<MODE>* **-theme** “*<ThemeName>*” (to delete granules associated with the specified theme and with retention priority less than or equal to the configured default limit and expiration date/time on or before midnight of the previous day). *Note:* The quotes around the theme name are required.
- **EcDlCleanupDataPool.pl** *<MODE>* **-themexref** “*<ThemeName>*” (to remove all granule cross references to a the specified theme from the Data Pool

inventory, without deleting the granules or the physical granule files.) **Note:** The quotes around the theme name are required.

- **EcDlCleanupDataPool.pl <MODE> -orphan -phantom** (to validate the Data Pool by checking for orphans and phantoms and removing any discrepancies found from the Data Pool inventory and disks).
 - **EcDlCleanupDataPool.pl <MODE> -orphan -phantom -nofix** (to validate the Data Pool by checking for orphans and phantoms and log any discrepancies found without removing the discrepant data or inventory entries).
 - **EcDlCleanupDataPool.pl <MODE> -orphan -phantom -collgroup “<NAME1>,<NAME2>,<NAME3>, . . . <NAMEn>”** (to initiate a Data Pool validation but limit the validation to specified collection group(s), where the collection groups are separated by commas in a string set off by quotation marks).
 - **EcDlCleanupDataPool.pl <MODE> -cleanvalidate -orphan -phantom** (to initiate a Data Pool cleanup followed by validation, first cleaning up all granules in the specified mode with retention priorities less than or equal to the configured default limit and expiration date/time on or before midnight of the previous day, and then checking for and removing orphans and phantoms).
- The Cleanup Utility runs and the Cleanup Utility log file **EcDlCleanup.log** records errors, warnings, and information about utility events.
-

The Data Pool Cleanup Utility may be run with **cron** to execute it on a daily basis at a consistent time of day. The following procedure provides an example of adding a line to a **crontab** file to execute the OPS mode Data Pool Cleanup at 1:00 a.m. every day.

Establish Data Pool Cleanup to Run with cron

- 1 Log in at an ECS platform using an account with privileges to remove science, metadata, and browse files from Data Pool disks.
- 2 To ensure that the **crontab** command launches the **vi** editor, type **setenv EDITOR vi** and then press the **Return/Enter** key.
 - It may be desirable to include this command in the operator's **.cshrc** file to set the **crontab** editor to **vi** as part of the environmental settings normally used routinely.
- 3 Type **crontab -e** and then press the **Return/Enter** key.
 - The contents of the file are displayed, and the cursor is displayed on the first character at the upper left corner of the file. **Note:** If the operator has no **crontab** file on the current platform, this command opens a new one for editing.

- 4 If necessary, use the down arrow key on the keyboard to move the cursor down to a blank line.
 - The cursor is displayed at the beginning of the selected line.
 - 5 Type **i** to put the **vi** editor into the insert mode.
 - The **vi** editor is in the insert mode, but no feedback is provided.
 - 6 Type **0 1 * * * /usr/ecs/OPS/CUSTOM/utilities/EcDIDataPoolCleanup.pl OPS**.
 - The typed entry appears to the left of the cursor.
 - 7 Press the **Esc** key.
 - The cursor moves one character to the left and the **vi** editor is in the command mode.
 - 8 Type **:wq** and then press the **Return/Enter** key.
 - UNIX displays a message identifying the number of lines and characters in the **crontab** file (stored in the directory **/var/spool/cron/crontabs**) and then displays the UNIX prompt.
-

Running the Data Pool Access Statistics Utility

The Data Pool Access Statistics Utility (DPASU) parses logs of the Data Pool Web Access service and the FTP access service and stores the results in tables in the Data Pool database. The DPASU is a command-line utility that permits an option of entering input parameters. It is intended to be run with **cron** to cover an arbitrary 24-hour period starting at a time specified as a configuration parameter in a configuration file. However, an operator may run the utility from the command line specifying a start date as an input parameter to cover a period other than the normal 24-hour period addressed by **cron** or to cover that normal period if **cron** failed to process the logs for that period.

There are two versions of the DPASU, one for each type of log processed. The script named **EcDIRollupWebLogs.pl** runs on the Data Pool Web Access server and processes its log; its configuration file is **EcDIRollupWebLogs.CFG**. The script named **EcDIRollupFtpLogs.pl** runs on a server with access to SYSLOG with FTP access entries; its configuration file is **EcDIRollupFtpLogs.CFG**. These scripts capture data on downloads from the Data Pool, including date and time of access, path and file name of the file, and size of the file. The captured data are written to a temporary "flat file" -- a tab-delimited text file -- stored in the directory **/<ECS_HOME>/<MODE>/CUSTOM/data/DPL/**. The flat file is then exported to Sybase and stored in a table. The DPASU calls Sybase stored procedures to generate a separate rollup table, removes the flat file, and enters a record in a separate table identifying which periods have been rolled up in order to prevent inadvertent reprocessing of that period.

To prevent potential table locking, **cron** runs of the DPASU scripts should be separated so that they are not both running concurrently (e.g., separate their start times by at least 20 minutes). Use the following procedure to specify a 1:00 a.m. start time for the rollup and add a line to the

crontab files to run the DPASU for the OPS mode beginning at 2:00 a.m. every day with a 20-minute separation between the scripts.

Specify Data Pool Access Statistics Rollup Start Time and DPASU Execution with *cron*

- 1 Log in at the host for EcDIRollupWebLogs.pl and its configuration file (e.g., e0mss21, g0mss21, l0mss21, n0mss21).
- 2 To change to the directory containing the configuration file, type the command **cd /usr/ecs/OPS/CUSTOM/cfg** and then press the **Return/Enter** key.
 - The working directory is changed to **/usr/ecs/OPS/CUSTOM/cfg**.
- 3 To look at the Rollup Start Time specified in the configuration file, type **vi EcDIRollupWebLogs.CFG** and then press the **Return/Enter** key.
 - The contents of the file are displayed, and the last line of the file indicates the start time in format similar to the following:

```
ROLLUP_START_TIME=3:00
```

and the cursor is displayed on the first character at the upper left corner of the file.
 - If the start time is correct, exit **vi** by typing **:q!** and pressing the **Return/Enter** key; then go to Step 10. Otherwise, to change the time, execute Steps 4 - 9.
- 4 Use the arrow keys on the keyboard to move the cursor down to the line specifying the **ROLLUP_START_TIME** and to move it to the right until it is located over the first character in the time value.
 - The cursor is moved to the start time location; the line should look similar to the following:

```
ROLLUP_START_TIME=3:00
```
- 5 Type **x** to delete the number under the cursor.
 - The number is deleted; the line should look similar to the following.

```
ROLLUP_START_TIME=:00
```
 - **Note:** If more characters in the time value are to be changed, you can type **x** repeatedly to delete additional characters. For this exercise, you need only delete one character.
- 6 Type **i** to put the **vi** editor into the insert mode.
 - The **vi** editor is in the insert mode, but no feedback is provided.
- 7 Type **1**.
 - The typed entry appears to the left of the cursor.
- 8 Press the **Esc** key.
 - The cursor moves one character to the left and the **vi** editor is in the command mode.

- 9 Type **ZZ** (be sure to use upper case).
 - The file is saved and the UNIX prompt is displayed.
- 10 To ensure that the **crontab** command launches the **vi** editor, type **setenv EDITOR vi** and then press the **Return/Enter** key.
 - It may be desirable to include this command in the operator's **.cshrc** file to set the **crontab** editor to **vi** as part of the environmental settings normally used routinely.
- 11 Type **crontab -e** and then press the **Return/Enter** key.
 - The contents of the file are displayed, and the cursor is displayed on the first character at the upper left corner of the file. *Note:* If the operator has no **crontab** file on the current platform, this command opens a new one for editing.
- 12 If necessary, use the down arrow key on the keyboard to move the cursor down to a blank line.
 - The cursor is displayed at the beginning of the selected line.
- 13 Type **i** to put the **vi** editor into the insert mode.
 - The **vi** editor is in the insert mode, but no feedback is provided.
- 14 Type **0 2 * * * /usr/ecs/OPS/CUSTOM/utilities/EcDIRollupWebLogs.pl OPS -noprompt**.
 - The typed entry appears to the left of the cursor.
- 15 Press the **Esc** key.
 - The cursor moves one character to the left and the **vi** editor is in the command mode.
- 16 Type **:wq** and then press the **Return/Enter** key.
 - UNIX displays a message identifying the number of lines and characters in the **crontab** file (stored in the directory **/var/spool/cron/crontabs**) and then displays the UNIX prompt.
- 17 Log in at the host for EcDIRollupFtpLogs.pl and its configuration file (e.g., e0dps01, g0dps01, l0dps01, n0dps01).
- 18 To change to the directory containing the configuration file, type the command **cd /usr/ecs/OPS/CUSTOM/cfg** and then press the **Return/Enter** key.
 - The working directory is changed to **/usr/ecs/OPS/CUSTOM/cfg**.

- 19 To look at the Rollup Start Time specified in the configuration file, type **vi EcDIRollupFtpLogs.CFG** and then press the **Return/Enter** key.
- The contents of the file are displayed, and the last line of the file indicates the start time in format similar to the following:

```
ROLLUP_START_TIME=3:00
```


and the cursor is displayed on the first character at the upper left corner of the file.
 - If the start time is correct, exit **vi** by typing **:q!** and pressing the **Return/Enter** key; then go to Step 21. Otherwise, to change the time, execute Step 20.
- 20 Repeat Steps 4-9 to change the time in **EcDIRollupFtpLogs.CFG**.
- 21 To ensure that the **crontab** command launches the **vi** editor, type **setenv EDITOR vi** and then press the **Return/Enter** key.
- It may be desirable to include this command in the operator's **.cshrc** file to set the **crontab** editor to **vi** as part of the environmental settings normally used routinely.
- 22 Type **crontab -e** and then press the **Return/Enter** key.
- The contents of the file are displayed, and the cursor is displayed on the first character at the upper left corner of the file. *Note:* If the operator has no **crontab** file on the current platform, this command opens a new one for editing.
- 23 If necessary, use the down arrow key on the keyboard to move the cursor down to a blank line.
- The cursor is displayed at the beginning of the selected line.
- 24 Type **i** to put the **vi** editor into the insert mode.
- The **vi** editor is in the insert mode, but no feedback is provided.
- 25 Type **20 2 * * * /usr/ecs/OPS/CUSTOM/utilities/EcDIRollupFtpLogs.pl OPS -noprompt**.
- The typed entry appears to the left of the cursor.
- 26 Press the **Esc** key.
- The cursor moves one character to the left and the **vi** editor is in the command mode.
- 27 Type **:wq** and then press the **Return/Enter** key.
- UNIX displays a message identifying the number of lines and characters in the **crontab** file (stored in the directory **/var/spool/cron/crontabs**) and then displays the UNIX prompt.
-

Although the Data Pool Access Statistics Utility scripts are intended to be run with **cron**, if it is necessary to run them from the command line, it is possible to do so. For example, if **cron** fails to complete successfully for any reason, no entry is made into the record table to indicate that a period was processed. In that event, the statistics can be captured for the missing interval by running the utility manually.

There are seven command-line parameters for use with the utility scripts (see 609 document information):

- The **<MODE>** parameter indicates the mode (must specify a valid directory path) in which the script is to run; it is mandatory, unlabeled, and must be the first parameter following the command.
- The **-noprompt** parameter optionally specifies suppression of output to the screen.
- The **-nodelete** parameter optionally prevents the flat file from being deleted upon completion of the run.
- The **-flatfile <path/file>** parameter optionally provides an alternative path/file name for the flat file produced by the parser (useful only with the **-nodelete** option).
- The **-ftp <path/file>** parameter optionally indicates an alternative ftp log path/file(s) to be used instead of the configured default path/file (for the **EcDIRollupFtpLogs.pl** script only). Wildcards may be used, but must be escaped (i.e., preceded with a \).
- The **-web <path/file>** parameter optionally indicates an alternative web log path/file(s) to be used instead of the configured default path/file (for the **EcDIRollupWebLogs.pl** script only). Wildcards may be used, but must be escaped (i.e., preceded with a \).
- The **-start <date>** parameter optionally indicates an alternative start date for the rollup period, using the format MM/DD, and may be used to process a previously uncovered period.

With the exception of the mandatory **<MODE>** parameter, which must appear first after the command, the other parameters may be used in various orders and combinations. For example, to run without screen prompts or information, starting from December 22, and to retain the flat file, the command for accumulating statistics on web access should be entered as follows:

EcDIRollupWebLogs.pl OPS -noprompt -nodelete -start 12/22.

To run with normal screen information display, starting from February 15, but using an alternative file with wildcards for the web log, the command should be similar to the following:

EcDIRollupWebLogs.pl OPS - start 2/15 -web /usr/var/*.log.

Use the following procedure to run the Data Pool Access Statistics Utility scripts from the command line, with normal screen information display.

- The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/utilities**.



The logo for the Data Pool Access Statistics Utility. It features a stylized 'D' composed of a grid of squares. The top-left square contains the number '0'. The top-right square contains a diagonal line from the bottom-left to the top-right. The bottom-left square contains a diagonal line from the top-left to the bottom-right. The bottom-right square contains a diagonal line from the top-left to the bottom-right. Below the grid, the text 'Data Pool Access Statistics Utility' is written in a sans-serif font.

The DPASU will examine the logs for access entries between the following times:

START:	11	26	03	00
END:	11	27	02	59

```
File list:
/usr/ecs/OPS/COTS/www/ns-home/www/logs/access
Processing Web logs...
No access entries found in any of the Web logs
```

```
Rollup successful!  
Removing flat file...OK  
Gracefully exiting...
```

- Log in at the host for EcDIIRollupFtpLogs.pl and its configuration file (e.g., e0dps01, g0dps01, l0dps01, n0dps01).

- The working directory is changed to `/usr/ecs/<MODE>/CUSTOM/utilities`.

[illegible]

The DPASU will examine the logs for access entries between the following times:

Checking for already covered rollup periods...

```
/var/adm/SYSLOG
```

Processing FTP logs...

No access entries found in any of the FTP logs

```
Cleaning up table "DlFtpAccessLog"...OK
```

Exporting flat file to Sybase...OK

No access data was available to roll up.
DPASU will skip this step.

Rollup successful!

Removing flat file...OK

Gracefully exiting...

The three remaining utilities are shell scripts for archiving, deleting, and restoring information in database tables populated by the DPASU. The **Data Pool Archive Access Statistics Data Utility** is run from the command line as needed or desirable to connect to the Data Pool database and write granule access data for a specified time range from the DIGranuleAccess, DIGranuleSubscription, and DIAccessRollup tables to an ASCII file. Once this is done, the operator can run the **Data Pool Delete Access Statistics Data Utility** from the command line to

delete the archived data from the Data Pool database. If it is desirable to restore deleted data to the database, the **Data Pool Restore Access Statistics Data Utility** can be run from the command line to restore the data. The following procedures are applicable.

Archive Access Statistics using the Data Pool Archive Access Statistics Data Utility

- 1 Log in at the host for the Data Pool database (e.g., e0acg11, g0acg01, l0acg02, n0acg01).
 - 2 To change directory to the directory containing the Data Pool Archive Access Statistics Data Utility, type **cd /usr/ecs/<MODE>/CUSTOM/dbms/DPL** and then press the **Return/Enter** key.
 - The working directory is changed to **cd /usr/ecs/<MODE>/CUSTOM/dbms/DPL**.
 - 3 Type **DIDbArchiveAccessStat <MODE> <STARTDATE> <STOPDATE> <ARCHIVEDIR> <USERNAME> <SERVER> <DBNAME>** and then press the **Return/Enter** key.
 - **Note:** <MODE> is the mode in which the utility is being executed (e.g., OPS, TS1, TS2). <STARTDATE> is the start date time range, in format **yyyymmdd**, for the data to be archived. <STOPDATE> is the stop date time range, in format **yyyymmdd**, for the data to be archived. <ARCHIVEDIR> is the absolute path where the generated ASCII files are to be stored. <USERNAME> is the Sybase login name. <SERVER> is the Sybase Server for the Data Pool database (e.g., e0acg11_srvr, g0acg01_srvr, l0acg02_srvr, n0acg01_srvr). <DBNAME> is the name of the Data Pool database (e.g., DataPool OPS).
 - The script displays a prompt for entry of the password for the Sybase login.
 - 4 Type <password> and then press the **Return/Enter** key (**Note:** This may require input from the Database Administrator).
 - The script runs and the Archive Access Statistics Utility log file **DIDbArchiveAccessStat.log** records errors, warnings, and information about utility events. The log is written to the directory **/usr/ecs/<MODE>/CUSTOM/logs**.
-

To run the Data Pool Delete Access Statistics Data Utility, use the following procedure.

Delete Access Statistics using the Data Pool Delete Access Statistics Data Utility

- 1 Log in at the host for the Data Pool database (e.g., e0acg11, g0acg01, l0acg02, n0acg01).
- 2 To change directory to the directory containing the Data Pool Delete Access Statistics Data Utility, type **cd /usr/ecs/<MODE>/CUSTOM/dbms/DPL** and then press the **Return/Enter** key.
 - The working directory is changed to **cd /usr/ecs/<MODE>/CUSTOM/dbms/DPL**.

- 3 Type **DIDbDeleteAccessStat** *<MODE>* *<STARTDATE>* *<STOPDATE>* *<USERNAME>* *<SERVER>* *<DBNAME>* and then press the **Return/Enter** key.
 - **Note:** *<MODE>* is the mode in which the utility is being executed (e.g., OPS, TS1, TS2). *<STARTDATE>* is the start date time range, in format *yyyymmdd*, for the data to be deleted. *<STOPDATE>* is the stop date time range, in format *yyyymmdd*, for the data to be deleted. *<USERNAME>* is the Sybase login name. *<SERVER>* is the Sybase Server for the Data Pool database (e.g., e0acg11_svr, g0acg01_svr, l0acg02_svr, n0acg01_svr). *<DBNAME>* is the name of the Data Pool database (e.g., DataPool_OPS).
 - The script displays a prompt for entry of the password for the Sybase login.
 - 4 Type *<password>* and then press the **Return/Enter** key (**Note:** This may require input from the Database Administrator).
 - The script runs and the Delete Access Statistics Utility log file **DIDbDeleteAccessStat.log** records errors, warnings, and information about utility events. The log is written to the directory */usr/ecs/<MODE>/CUSTOM/logs*.
-

To run the Data Pool Restore Access Statistics Data Utility, use the following procedure.

Restore Access Statistics using the Data Pool Restore Access Statistics Data Utility

- 1 Log in at the host for the Data Pool database (e.g., e0acg11, g0acg01, l0acg02, n0acg01).
- 2 To change directory to the directory containing the Data Pool Restore Access Statistics Data Utility, type **cd /usr/ecs/<MODE>/CUSTOM/dbms/DPL** and then press the **Return/Enter** key.
 - The working directory is changed to **cd /usr/ecs/<MODE>/CUSTOM/dbms/DPL**.
- 3 Type **DIDbRestoreAccessStat** *<MODE>* *<STARTDATE>* *<STOPDATE>* *<ARCHIVEDIR>* *<USERNAME>* *<SERVER>* *<DBNAME>* and then press the **Return/Enter** key.
 - **Note:** *<MODE>* is the mode in which the utility is being executed (e.g., OPS, TS1, TS2). *<STARTDATE>* is the start date time range, in format *yyyymmdd*, for the data to be restored. *<STOPDATE>* is the stop date time range, in format *yyyymmdd*, for the data to be restored. *<ARCHIVEDIR>* is the absolute path of the storage location for the ASCII files containing the data to be restored. *<USERNAME>* is the Sybase login name. *<SERVER>* is the Sybase Server for the Data Pool database (e.g., e0acg11_svr, g0acg01_svr, l0acg02_svr, n0acg01_svr). *<DBNAME>* is the name of the Data Pool database (e.g., DataPool_OPS).
 - The script displays a prompt for entry of the password for the Sybase login.

- 4 Type *<password>* and then press the **Return/Enter** key (*Note*: This may require input from the Database Administrator).
 - The script runs and the Archive Access Statistics Utility log file **DIDbRestoreAccessStat.log** records errors, warnings, and information about utility events. The log is written to the directory **/usr/ecs/<MODE>/CUSTOM/logs**.

Using the Batch Insert Utility

The Batch Insert Utility allows operators to specify Data Pool insert for granules residing in the ECS archive, as well as data from outside ECS (non-ECS granules). The utility queues the granules up for dispatch by the Data Pool Action Dispatcher (DPAD) for insertion by the Data Pool Insert Utility (DPIU). It accepts either a list of ECS granule identifiers or a list of non-ECS names; the list can be provided either as an input file or as standard input. A label identifying a batch of granules is specified as a command-line parameter, using the option **-label**, so that operators can monitor a batch with the DPM GUI.

Granules to be inserted can also be linked to a theme, using the option **-theme**. In fact, the Batch Insert Utility can also be used with that option to link granules already present in the Data Pool to a theme, or to additional themes. However, it is important to note that if the granules were originally inserted into the Data Pool using the Batch Insert Utility, you must use a different batch label when linking the granules to the theme than was used for the original insert. This is necessary because the Batch Insert Utility is designed to reject inserts that are in a batch with a label identical to one for which granules are already being processed. So, even if the batch has been inserted, if the inserts are still in the queue (e.g., with a status of **Completed**), you cannot run another batch with the same label to link them to a theme.

The following procedure is applicable.

Batch Insert of Data into the Data Pool

- 1 Log in at the machine where the Data Pool Batch Insert Utility is installed (e.g., e0dps01, g0dps01, l0dps01, n0dps01).
 - *Note*: The login must be as either cmshared or allmode to ensure correct permissions.
- 2 To change to the directory for starting the Batch Insert Utility, type **cd /usr/ecs/<MODE>/CUSTOM/utilities** and then press the **Return/Enter** key.
 - The working directory is changed to **/usr/ecs/<MODE>/CUSTOM/utilities**.

- 3 At the UNIX prompt, enter the command to start the Batch Insert Utility, in the form **EcDiBatchInsert.pl <MODE> -ecs | -nonecs [-file <pathname>] [options]**.
- **Note:** The following are examples of valid command-line entries for initiating the Batch Insert Utility:
 - **EcDiBatchInsert.pl <MODE> -ecs -file /home/cmshared/<filename>** (to add actions to the action insert queue for all ECS granules specified by granule IDs in the specified file. Because the command does not specify a **-label** parameter, the label is formed from the first 16 characters of the input file name).
 - **EcDiBatchInsert.pl <MODE> -nonecs -file /home/cmshared/<filename> -label Chig_volcano -theme “Chiginagak Volcano 2002”** (to add actions to the insert action queue for all non-ECS granules specified by XML pathnames in the specified input file, with all granules linked with the theme name “Chiginagak Volcano 2002” in the Data Pool database). **Note:** The theme name must already be in the Data Pool database in the DIThemes table; if necessary, use the DPM GUI **Manage Themes** tab to define the theme before running the batch insert.
 - **Note:** You can use Batch Insert with the **-theme** option to link granules already in the Data Pool to a theme, but if the granules were originally inserted using the Batch Insert Utility, you must use a different batch label than was used for the original insert; otherwise, the insert of the theme links may be rejected.
 - **EcDiBatchInsert.pl <MODE> -ecs -file /home/cmshared/<filename> -mdonly** (to add actions to the action insert queue for all ECS granules specified by granule IDs in the specified file, but insert metadata only. Because the command does not specify a **-label** parameter, the label is formed from the first 16 characters of the input file name).
 - **EcDiBatchInsert.pl <MODE> -ecs -file /home/cmshared/<filename> -rpriority 255** (to add actions to the action insert queue for all ECS granules specified by granule IDs in the specified file, and to set their retention priority to 255. Because the command does not specify a **-label** parameter, the label is formed from the first 16 characters of the input file name).
 - **EcDiBatchInsert.pl <MODE> -ecs -file /home/cmshared/<filename> -rpriority 255 -rperiod 10 -dpriority 5** to add actions to the action insert queue for all ECS granules specified by granule IDs in the specified file, and to set their retention priority to 255 and their retention period to 10 days, with dispatch priority set to 5. Because the command does not specify a **-label** parameter, the label is formed from the first 16 characters of the input file name).
 - The Batch Insert Utility runs and events and errors are recorded in the Batch Insert Utility log file **EcDiBatchInsert.log**.
-

Working with the DataPool Order Status & Control GUI

The **DataPool Order Status & Control** application is a set of HTML pages that allows the operator to view the status of orders and order items for Data Pool orders (*i.e.*, orders placed using the Data Pool Web Access GUI Shopping Cart order capability or the single granule converter dialog). It also allows the operator to control key aspects of the order process such as the queue control. The application is split into four functional areas: **Queue Control**, **Orders**, **Order Items**, and **Help**. Each page of the DataPool Order Status & Control application provides access to these four functions through links at the top of the page. The **help** function is the primary source of information in this section of the lesson and is therefore not addressed separately.

Orders Page

Figure 59 shows the **Orders** page of the DataPool Order Status & Control application, along with the **Order Details Viewer**. This page reports the status of orders and is primarily comprised of the order table.

The screenshot displays two web browser windows. The main window, titled "DataPool Order Status - Netscape", shows the "Orders" page. It features a header with navigation links: "Queue Control", "Orders", "Order Items", and "Help". Below the header, a summary states: "There are totally 23 orders, 1 order is currently being processed. There are also 2 failed orders which need operator attention. The filtering status in this page is: No filtering". A table lists the orders with columns: Order ID, Status, E-Mail, User Name, Creation Date, and Last Update. A mouse cursor points to the first row of the table. The table shows 10 rows of data, with the last row being a "FAILPERN" status. The right window, titled "DataPool Order Details Viewer - Netscape", shows the "Order Details Viewer" for Order ID 10. It displays details such as Status: DONE, Creation Date: 04/15/03, Last Update: 05/15/03, Email: sjsk@nsidc.org, and Real Name: s khalsa. It also shows coordinates and projection information.

Order ID	Status	E-Mail	User Name	Creation Date	Last Update
32	DONE	hconover@utsc.uah.edu	-	05/15/03	05/15/03
9	DONE	sjsk@nsidc.org	s khalsa	04/15/03	05/15/03
17	DONE	sjsk@nsidc.org	s khalsa	04/18/03	05/15/03
10	DONE	sjsk@nsidc.org	s khalsa	04/15/03	05/15/03
19	DONE	sjsk@nsidc.org	s khalsa	04/21/03	05/15/03
31	FAILPERN	btmclean@nsidc.org	-	05/14/03	05/15/03
30	DONE	btmclean@nsidc.org	-	05/14/03	05/14/03
29	DONE	marilyn@nsidc.org	-	05/12/03	05/12/03
27	FAILPERN	btmclean@nsidc.org	-	05/02/03	05/06/03
28	FAILPERN	btmclean@nsidc.org	-	05/02/03	05/02/03

Figure 59. DP Order Status & Control: Orders Page and Order Details

As the figure shows, the table rows each represent an order. The table header and footer contain controls for paging and filtering. The order table columns are:

- **Order ID** - an id representing the order. The id is automatically assigned by the database when an order is first entered. To navigate to the order items associated with the order, simply click on the order id link. This column also contains an icon link (🔗) to the order details page (**Order Details Viewer**). The order details page includes the fields shown in the order table plus additional details such as subsetting parameters and format and projection override settings. The order details page also has navigation links (first|previous|next|last) allowing step-wise navigation from one order to the next.
- **Status** - processing status of the order. Values are:
 - Entered - order is entered but not yet selected from the queue for processing.
 - Processing - order has been selected from the queue and at least one of its order items is being processed.
 - Processed - all order items have been processed and the order is not yet in the packaging phase.
 - Packaging - the order is complete and is being packaged. Packaging is essentially the processing of making sure all the necessary files are linked into the download directory.
 - FAILOPR - Order has failed processing and operator intervention is required, but the operator has not been notified (e-mail has not been sent).
 - FAILOPERN - Order has failed processing, operator intervention is required and notification has been sent to the operator e-mail address. When in this state, this column also contains two icon links: retry (🔄) and "mark as complete" (📁). Selecting re-try attempts to re-process all of the failed order items for this order. Selecting "mark as complete" leaves all failed order items as failed and sets the order to DONE (if there are successfully completed order items) or FAILED if all of the order items failed. In either case, an order notification e-mail is sent to the user once the order is complete or has been finally failed.
 - FAILED - order is failed on arrival and cannot be revived.
 - DONE - order is finished processing. Order items can be DONE, or, if order intervention was required, FAILED when the order is in this state. Once an order is in the DONE state, an order notification e-mail is shipped to the user.
- **E-Mail** - e-mail address of the user.
- **User Name** - User's first and last name.
- **Creation Date** - date the order was entered into the system.
- **Last Update** - last time the order was acted upon by the system.

- **Subsetting** - NOTE: subsetting parameters do not appear in the table, but can be viewed by clicking the order details icon (🔍).

The Order table header and footer consist of:

- **Filter** control/indicators:
 - On/Off - Turn filtering on or off. The default state is off.
 - By Status - only display orders of the selected status. The **Do Filtering** button must be activated to implement the filter settings.
- Column Configuration (⇄) - a click on this icon link displays a page which allows the changing of column order in the order table (the first two columns are fixed).
- **#Rows in Page** - changes the number of rows displayed on a page. Values are 10,20,30,40,and 50.
- **Showing** - allows arbitrary selection any of the current pages for display =.
- Navigation Controls - allows the step-wise selection of order pages.

Order Items Page

Figure 60 shows the **Order Items** page of the DataPool Order Status & Control GUI, along with the **Order Item Details Viewer**. The Order Items (Items) page reports the status of order items and is primarily comprised of the order item table. The table rows each represent an order item. The table header and footer contain controls for paging and filtering.

The Order Item table columns are:

- **Item ID** - an id representing the order item. The id is automatically assigned by the database when an order is first entered. This column also contains a link (🔍) to the order item details page (**Order Item Details Viewer**). This details page includes the fields shown in the order item table plus additional details such as the projection and format settings. It also allows step-wise navigation (first|previous|next|last) from order item to item. Note that this user interface allows navigation through all of the order items, not just the items in the currently selected order.
- **Status** - processing status of the order item. Values are:
 - <NULL> - order item is entered but is not yet being processed.
 - PROCESSING - order item is being processed.
 - FAILED - order item processing failed. Check the error code. When in this state, the column also contains a retry icon (🔄) which allows the operator to run the processing of the order item again.
 - DONE - order item has been processed.

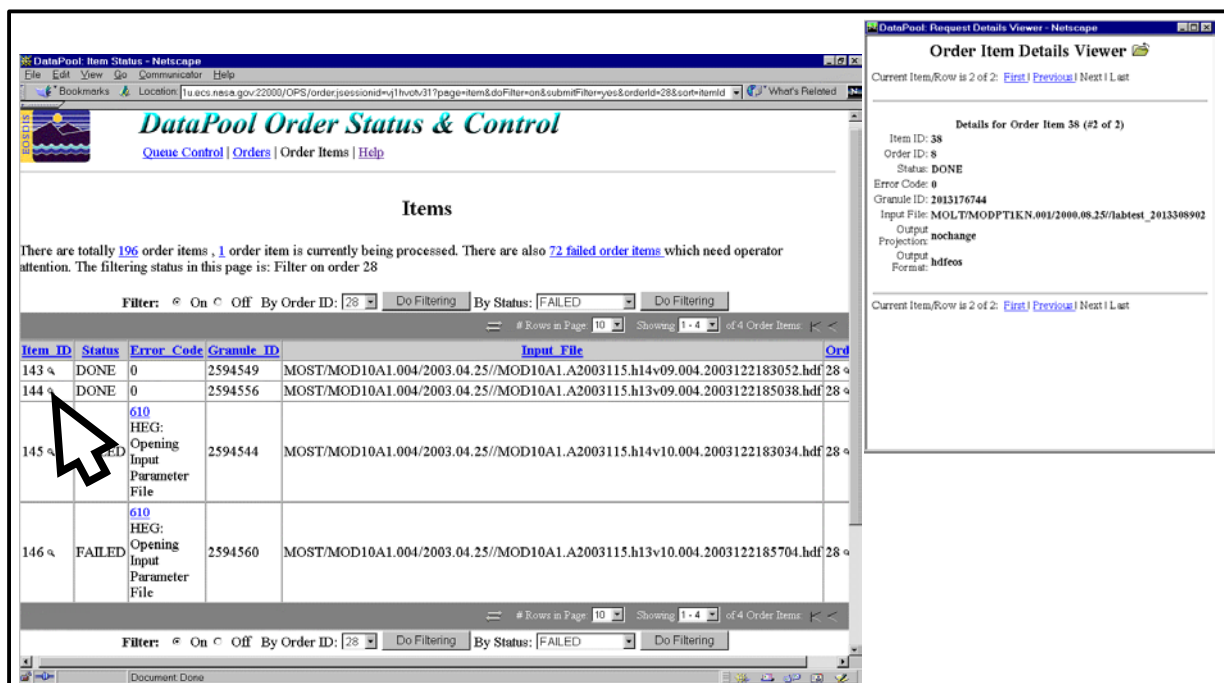


Figure 60. DP Order Status & Control: Order Items Page and Order Item Details

- **Error Code** - number indicating a failure reason. This code usually has a small bit of explanatory text after it. More details are available by clicking on the code link itself.
- **Granule ID** - the DataPool id for the granule. This is not the same id that is in the Science Data Server database.
- **Input File** - name of the input file (typically an HDF-EOS file since that is the only kind the HEG Converter currently supports).
- **Order ID** - an id represent the order. Click on the order link to go back to the order table.

The Order Item table header and footer consist of:

- **Filter** control/indicators:
 - On/Off - Turn filtering on or off. The default state is off.
 - By Order ID - shows only the items associated with the selected order id. This is the default when the order link is selected from the order page.
 - By Status - only display orders of the selected status. The **Do Filtering** button must be activated to implement the filter settings.
- Column Configuration - (☰) - this icon link displays a page which allows the changing of column order in the order item table (the first three columns are fixed).
- **#Rows in Page** - changes the number of rows displayed on a page. Values are 10,20,30,40,and 50.

- **Showing** - allows arbitrary selection any of the current pages for display.
- **Navigation Controls** - allows the step-wise selection of order pages.

Queue Control Page

Figure 61 shows the **Queue Control** page of the DataPool Order Status & Control GUI. The Queue Control page (Processing and Queue Status) allows the operator to start and stop the HEG (HDF-EOS to GeoTIFF Converter) Front End server and to configure the HEG-related order parameters. The HEG Front End Server is responsible for processing DataPool requests for data conversion. These requests are termed 'orders' because they are asynchronous in nature and require a management process very similar to ECS orders.

There are four controls within the Queue Control page:

- **HEG Front End Server** - there are three elements to this control:
 - **HEG Front End Server** Status Indicator: 'UP' or 'DOWN'. If the server is down, a **Start Up** button appears to the right of the status indicator. When the status is 'UP' both exit buttons (**Exit Gracefully** and **Exit Immediately (No Cleanup)**) appear to the right of the status indicator.
 - **Exit Gracefully** button - requests that the HEG Front End (a java process) exit, but leave existing conversion processes running.
 - **Exit Immediately (No Cleanup)** button - stops both the HEG Front End and all of its associated converter processes.

NOTE: These buttons do not require a click on the **Apply Changes** button (they are submit buttons and as such send requests directly).

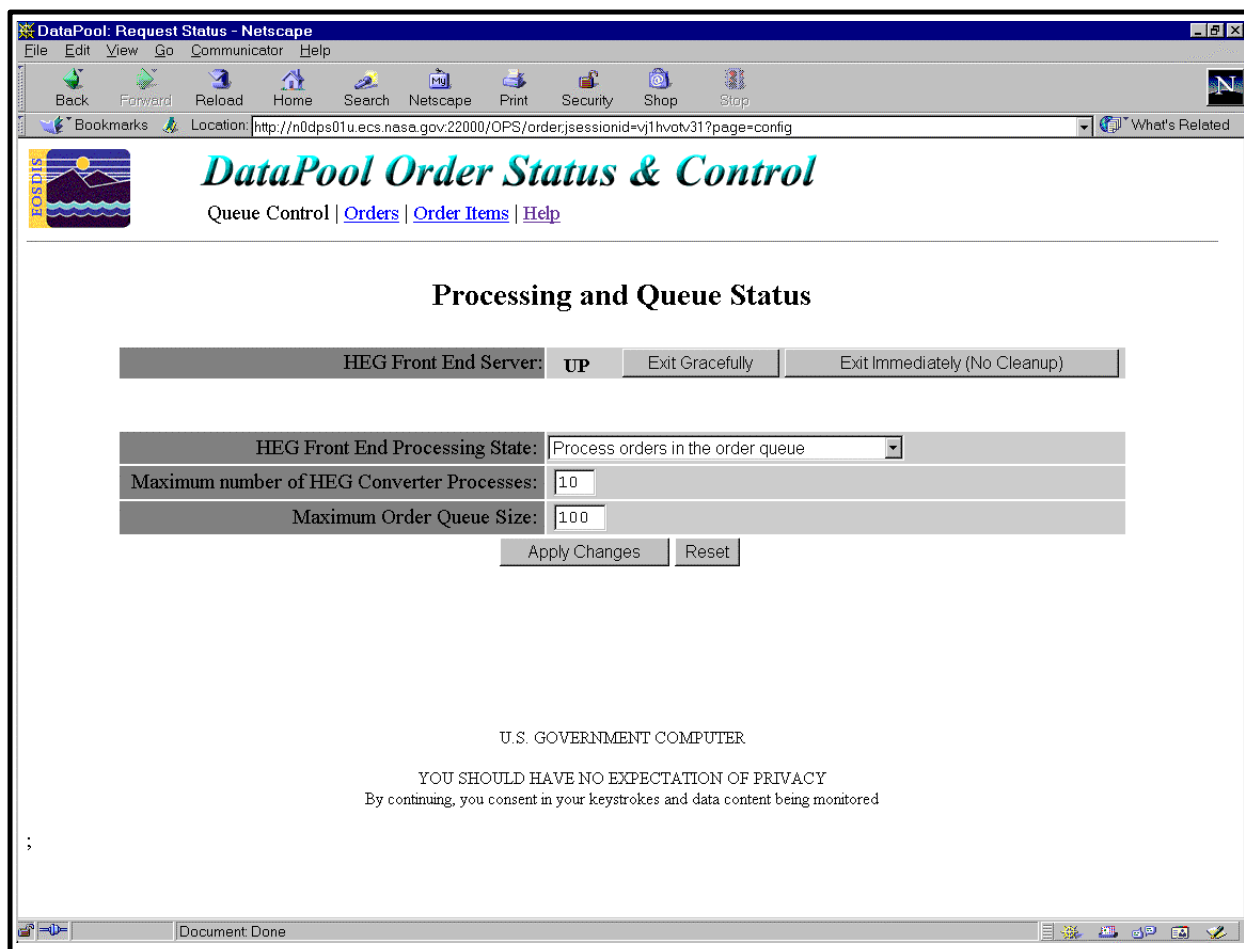


Figure 61. DataPool Order Status & Control: Queue Control Page

- **HEG Front End Processing State** - there are two processing states selectable from a pull-down list:
 - 'Process orders in the order queue' - this means that the HEG Front End will always try to process orders that are in the queue. The HEG Front End normally comes up in this state.
 - 'STOP processing orders in the order queue' - this means that the HEG Front End will not process any orders in the queue even though the front end itself will be running. This is useful if you decide that there are too many orders being processed and you want to wait until all de-queued orders are processed.
- **Maximum number of HEG Converter Processes** - this parameter limits the number of converter processes (each of these is a C executable that can potentially require significant computer resources to run). Once the limit is reached, orders wait for converter slots to become available. Individual resource requirements vary and are dependent on the format and type of projection requested.

- **Maximum Order Queue Size** - this parameter limits the size of the order queue. Once the limit is reached, user requests to place orders are refused.

Unless noted, each of these controls requires a click on the **Apply Changes** button before the settings take effect. The **Reset** button simply clears the form values.

Use the following procedure to launch the DataPool Order Status & Control GUI.

Launch the DataPool Order Status & Control GUI



- 1 At the UNIX command shell prompt, type **setenv DISPLAY *clientname*:0.0** and then press the **Return/Enter** key.
 - For *clientname*, use either the local terminal/workstation IP address or its machine name.
- 2 Start the log-in to a Netscape host by typing **/tools/bin/ssh *hostname*** (e.g., g0ins02, e0ins02, l0ins02, n0ins02) at the UNIX command shell prompt, and press the **Return/Enter** key.
 - If you receive the message, **Host key not found from the list of known hosts. Are you sure you want to continue connecting (yes/no)?** type yes (“y” alone does not work).
 - If you have previously set up a secure shell passphrase and executed **sshremote**, a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears; continue with Step 3.
 - If you have not previously set up a secure shell passphrase; go to Step 4.
- 3 If a prompt to **Enter passphrase for RSA key '<user@localhost>'** appears, type your **Passphrase** and then press the **Return/Enter** key. Go to Step 5.
- 4 At the **<user@remotehost>'s password:** prompt, type your **Password** and then press the **Return/Enter** key.
 - You are logged in and a UNIX command shell prompt is displayed.
- 5 Type **netscape** and then press the **Return/Enter** key.
 - The Netscape web browser is displayed.
- 6 Click in the **Netsite:** field.
 - The field is highlighted.

- 7 Type the Universal Resource Locator (URL) for the DataPool Order Status & Control GUI and then press the **Return/Enter** key.
 - The DataPool Order Status & Control GUI **Orders** page is displayed, offering links to access Data Pool order status and control functions (**Queue Control**, **Orders**, **Order Items**, and **Help**) and a table of information on orders currently in the order queue.

Use the following procedure to review the status of Data Pool orders and examine the order items in a particular order.

Use the DataPool Order Status & Control GUI to Review Orders and Order Items

- 1 Launch the DataPool Order Status & Control GUI (refer to procedure **Launch the DataPool Order Status & Control GUI**, page 237).
 - The DataPool Order Status & Control GUI **Orders** page is displayed, offering links to access Data Pool order status and control functions (**Queue Control**, **Orders**, **Order Items**, and **Help**) as well as a table of information on orders currently in the order queue.
 - You can observe an order of interest by locating it in the **Order ID** column, scrolling if necessary.
 - Filtering and sorting capabilities are available through controls in the table header and footer.
- 2 It may be useful to filter the list. For example, if you receive notice that an order failed and you want to restrict the listed orders to display only those that have failed and for which notice has been sent to you, click on the pull-down arrow at the end of the **By Status** field, select FAILOPERN, and then click on the **Do Filtering** button.
 - The table displays only failed orders for which a notice has been sent to the operator.
- 3 It may be helpful to sort the list, using the column-heading links. For example, if the window displays many failed orders and you are looking for a failed order from a specific user, click on the **User_Name** link at the top of the column listing user names.
 - The listed orders are sorted by user name, permitting the operator to see all the failed orders from the user of interest in a single contiguous block.
- 4 If desired, in the **Order_ID** column click on the magnifier icon (🔍) for the order of interest to obtain the **Order Details Viewer**.
 - The **Order Details Viewer** is displayed with additional details concerning the order of interest.
- 5 To close the **Order Details Viewer**, click on the folder icon (📁) in the viewer.
 - The **Order Details Viewer** is closed.

- 6 To check the status of items in the order of interest, click on the order ID link in the **Order ID** column.
 - The system displays the **Order Items** page with the **Items** table listing the items in the order for which the Order ID was clicked.
 - For each listed item, the **Items** table shows the **Item_ID**, **Status**, any **Error Code**, the **Granule_ID**, the **Input_File**, and the **Order_ID**.
 - 7 If desired, click on the magnifier icon () to obtain the **Order Item Details Viewer**.
 - The **Order Item Details Viewer** is displayed with additional details concerning the selected order item.
 - 8 To close the **Order Item Details Viewer**, click on the folder icon () in the viewer.
 - The **Order Item Details Viewer** is closed.
-





As noted previously, the DataPool Order Status & Control GUI provides icons permitting the operator to take action on orders and/or order items that have failed in a way that may allow an operator intervention to reprocessing of the items that failed. The system is designed to send the operator an e-mail notification of the necessity for intervention. The operator may intervene to re-try the order or order items, or may choose to mark the order in question as complete, which results in the order being placed in one of the following states:

- DONE – some items are successfully completed and any failed items are left in the failed state.
- FAILED – all items were failed and are left in the failed state.

The following procedure addresses operator intervention in the case of a failed Data Pool order.

Intervene in a Failed Data Pool Order Susceptible to Operator Intervention

- 1 Launch the DataPool Order Status & Control GUI (refer to procedure **Launch the DataPool Order Status & Control GUI**, page 237).
 - The DataPool Order Status & Control GUI **Orders** page is displayed, offering links to access Data Pool order status and control functions (**Queue Control**, **Orders**, **Order Items**, and **Help**) as well as a table of information on orders currently in the order queue.
 - You can observe an order of interest by locating it in the **Order ID** column, scrolling if necessary.
 - Filtering and sorting capabilities are available through controls in the table header and footer.

- 2 Locate the order of interest (e.g., an order for which you receive e-mail notification of failure that may permit intervention to recover).
 - See procedure **Use the DataPool Order Status & Control GUI to Review Orders and Order Items**, steps 2 and 3 (page 229).
 - The **Status** column for a failed order of which the operator is notified indicates **FAILOPERN** and also contains two icon links: retry () and "mark as complete" ().
- 3 If desirable, review detailed information on the order and its order items (see the procedure **Use the DataPool Order Status & Control GUI to Review Orders and Order Items**, steps 4 - 8 (page 229).
- 4 To retry a failed order, click on the retry icon ().
 - A retry confirmation message asks **You have opted to retry an order. This will automatically retry all of the failed items within the order if there are any. Are you sure you want to do this?** and offers **Yes** and **No** buttons.
- 5 Click on the **Yes** button in the retry confirmation message.
 - A confirmation message indicates **Your request has been submitted. Please reload the corresponding page.** and offers a **Reload** button.
- 6 Click on the **Reload** button in the confirmation message.
 - On the **Orders** page, the **Status** of the order is shown as **ENTERED**.
 - On the **Order Items** page, the **Status** of the items in the order should be **NULL** (entered but not yet being processed) or **PROCESSING**.
- 7 To mark a failed order as complete, click on the "Mark as Complete" icon ().
 - A "Mark as Complete" confirmation message asks **You have opted to complete this order despite its failure status. This means that at least some granules in this order will not be delivered to the user as requested. Marking this order for completion means that it will never be able to be retried (at least from this user interface). Note that the user e-mail notice will automatically be sent at this point. If you wish to enter comments FOR THE USER regarding this order, please enter them below.** (A text entry box is provided for the entry.) **Change order status and send e-mail?** The message then offers **Yes** and **No** buttons.
- 8 Click on the **Yes** button in the "Mark as Complete" confirmation message.
 - A confirmation message indicates **Your request has been submitted. Please reload the corresponding page.** and offers a **Reload** button.

- 9 Click on the **Reload** button in the confirmation message.
 - On the **Orders** page, the **Status** of the order is shown as **DONE**.
 - On the **Order Items** page, the **Status** of the items in the order is shown as **FAILED**.
-

The DataPool Order Status & Control GUI **Queue Control** function provides a means for starting and stopping the HEG Front End Server (there is also a script named **EcDIHEGFrontEndControl** on the Data Pool host that provides an alternative method of starting and stopping the HEG Front End). The Queue Control function of the GUI also permits an operator to set the HEG Front End Processing State to process orders in the converter order queue or not, and to set limits on the maximum number of HEG Converter processes and the maximum number of orders in the Order Queue. The following procedure is applicable.

Use DataPool Order Status & Control GUI to Manage HEG Converter Front End Server

- 1 Launch the DataPool Order Status & Control GUI (refer to procedure **Launch the DataPool Order Status & Control GUI**, page 237).
 - The DataPool Order Status & Control GUI **Orders** page is displayed, offering links to access Data Pool order status and control functions (**Queue Control**, **Orders**, **Order Items**, and **Help**) as well as a table of information on orders currently in the order queue.
- 2 Click on the **Queue Control** link near the top of the display.
 - The **Processing and Queue Status** page is displayed.
- 3 To check the status of the HEG Front End Processing Server, observe the **HEG Front End Processing Server** line on the display.
 - The appearance of the line provides indication of the status, as follows:
 - If the server is up, the line indicates **UP** and there are two control buttons one labeled **Exit Gracefully** and one labeled **Exit Immediately (No Cleanup)**.
 - If the server is down, the line indicates **DOWN** and there is one control button labeled **Start Up**.
 - If the server is down and you wish to start it, go to Step 6.
 - If the server is up and you wish to stop it, go to Step 4 (graceful exit) or 5 (immediate exit); otherwise continue with Step 7.
- 4 To stop the server gracefully, click on the **Exit Gracefully** button.
 - The server does not begin any new processes and, after the server completes ongoing converter processes and exits, the **HEG Front End Processing Server** line on the display indicates **DOWN** and offers a **Start Up** button.

- 5 To stop the server immediately, click on the **Exit Immediately (No Cleanup)** button.
 - All ongoing converter processes are killed and the server exits immediately; the **HEG Front End Processing Server** line on the display indicates **DOWN** and offers a **Start Up** button.
 - 6 To start the server, click on the **Start Up** button.
 - The server starts; the **HEG Front End Processing Server** line on the display indicates **Up** and offers an **Exit Gracefully** button and an **Exit Immediately (No Cleanup)** button.
 - 7 If you wish to change the other parameters available on the **Processing and Queue Status** page, enter the new value in the input field, using the appropriate editing method:
 - To change the **HEG Front End Processing State**, click on the pull-down arrow at the end of the field and select the desired state (**Process orders in the order queue** or **STOP processing orders in the order queue**).
 - To change the value for the **Maximum number of HEG Converter Processes** or **Maximum Order Queue Size**, click at the end of the field and either use the **Backspace** key to delete the current value or use the mouse with the primary key held down to drag the cursor and highlight the current value. Then type the desired new value.
 - The field(s) display the desired new value(s).
 - 8 Click on the **Apply Changes** button.
 - The screen is refreshed with the new value(s) implemented.
 - The **Reset** button may be used to clear the form values.
-

Practical Exercises

Introduction

This exercise is designed to practice key elements of the Archive procedures. Perform the tasks identified in the exercise.

Equipment and Materials

One ECS workstation, the STK Powderhorn storage facility, a copy of 609-CD-610-003 *Release 6B Operations Tools Manual*, and a copy of 611-CD-610-002 *Mission Operation Procedures for the ECS Project*.

Perform Activities Related to Archive Processing

1. Locate the STK Powderhorn storage facility and the AMASS host. Point out the elements and sequence involved in starting AMASS.
2. At the STK Powderhorn, locate the control panels necessary for power up (or down) and identify all Power Switches.
3. Following all safety precautions, vary the STK Powderhorn offline and enter the unit; leave the unit and restore it to online status.
4. Launch the Data Distribution Graphical User Interface (GUI); examine the list of distribution requests. Then filter the list to examine only those requests that are staging.
5. Insert several granules that can be used to exercise the granule deletion capability, and then use that capability to delete the granules from the inventory and archive, specifying a lag time other than 0 in the Deletion Cleanup Utility script.
6. Use automatic loading procedures to load a 9940 tape into the STK Powderhorn. Then use automatic unloading procedures to remove the media you just loaded.
7. Experiment with the **vollist**, **dirfilelist**, and **volfilelist** commands for AMASS. Describe for yourself how the commands can be used to help you manage the archive.
8. Use the **vgexport -q** command to create a backup for the AMASS database.
9. Check for the existence of required running AMASS daemons; run **healthcheck** to determine the health of basis AMASS functions.
10. Look at the AMASS queue by using the **quedisplay** command.
11. Run the **amass_log** script to display AMASS messages in **/var/adm/SYSLOG** system log file (on SGI machines) or in **/var/adm/messages** (on Sun machines).
12. Launch the Data Pool Maintenance (DPM) GUI and review the Data Pool active insert processes. Use the **Refresh** button to obtain an immediate screen refresh. Change the automatic screen refresh rate and observe that the new rate takes effect.

13. Use the DPM GUI to suspend Data Pool Insert actions. Then use the GUI to resume Data Pool Insert actions.
14. Use the DPM GUI to check the Data Pool Insert Queue.
15. Use the DPM GUI **Manage Configuration Parameters** tab and increase the number of allowed active insert processes by 5. Then decrease the number by 5. Change the Default Retention Priority to a number of your choosing.
16. Use the DPM GUI to display a list of collection groups for data in the Data Pool. Select one of the groups and display the list of its collections. Then pick one of the collections in the group and display a description for it.
17. Use the DPM GUI to obtain a list and descriptions of themes in the Data Pool database. Then add a theme which is insert enabled but not web visible. Delete the theme you created.
18. Launch the Spatial Subscription Server (NSBRV) GUI and display a list of ECS events for which a subscription can be created.
19. Using the NSBRV GUI, display subscriptions in the NSBRV database and filter the list to show only subscriptions for data with a selected short name. Filter the list again to display only subscriptions for data with a different short name.
20. Use the NSBRV GUI to obtain a list of bundling orders and view the details of one of them. then add a bundling order, but do not associated it with a subscription. Delete the bundling order you created.
21. Use the NSBRV Command Line Interface to view the contents of a subscription in the NSBRV database.
22. Use the NSBRV GUI to view any acquire and notification actions in the Action Queue.
23. Use the NSBRV GUI to display statistics on NSBRV processing of events and actions.
24. Run the **Update Granule Utility** to extend the period of retention for a single granule in the Data Pool. Examine the log file for the utility and review the entries for the update.
25. Invoke the **Data Pool Cleanup Utility** from the command line.
26. Execute the **Data Pool Access Statistics Utility** from the command line.
27. Use the **Data Pool Archive Access Statistics Data Utility** to archive statistics on Data Pool access for the last seven days.
28. Use the **Batch Insert Utility** to insert several granules from the archive into the Data Pool.
29. Launch the DataPool Order Status & Control GUI and review Data Pool orders and their order items.

Slide Presentation

Slide Presentation Description

The following slide presentation represents the slides used by the instructor during the conduct of this lesson.

This page intentionally left blank.